

# **Project Execution Plan**

Final Version

## **FutureGrid: An Experimental, High-Performance Grid Test-Bed**

Prepared by  
Indiana University  
for the  
National Science Foundation

**5 August 2009**

## Table of Contents

|  |    |
|--|----|
| 1. Executive Summary .....   | 5  |
| 2. Science Plan .....  | 6  |
| 2.1 Motivation and Purpose .....   | 6  |
| 2.2 Early Science Experiences .....  | 8  |
| 2.3 Attracting and Selecting Interesting and Valuable Research.....            | 10 |
| 2.4 Detailing Capabilities of FutureGrid .....                                 | 11 |
| 2.5 Identifying Research Problems and Allocating Usage .....                   | 12 |
| 3. Organizational Roles.....   | 12 |
| 3.1 Funded University Participants.....  | 12 |
| 3.2 Unfunded University Participants.....                                      | 17 |
| 3.3 Private Sector Partners .....  | 17 |
| 4. Management Plan.....  | 17 |
| 4.1 Overall Management Structure .....   | 17 |
| 4.2 Key Management Personnel .....   | 18 |
| 4.3 Committee Structures.....  | 19 |
| 4.4 Consensus Management Process .....   | 20 |
| 4.5 Maintaining, Refreshing, and Executing the Project Vision.....             | 20 |
| 5. Project Deliverables .....  | 20 |
| 5.1 Facility .....   | 20 |
| 5.2 Software .....   | 21 |
| 5.3 Educational Materials .....  | 22 |
| 5.4 Scientific Data and Knowledge .....  | 22 |
| 5.5 Better Educated Students .....   | 22 |
| 5.6 Reports, Presentations, and Published Works.....                           | 22 |
| 6. Minority Serving Institution Engagement Plan .....                          | 23 |
| 6.1 Goals of Engagement.....   | 23 |
| 6.2 Leverage of MSI Cyberinfrastructure Empowerment Coalition (MSI-CIEC) ..... | 24 |
| 6.3 Leveraging Social Networking Technologies.....                             | 25 |
| 6.4 Workshops and Tutorials Specific to MSI Engagement.....                    | 25 |
| 7. Project Budget and Work Breakdown Structure.....                            | 25 |
| 7.1 Project Budget.....  | 26 |
| 7.2 Methodology and Assumptions Used for Estimating Budget Components .....    | 26 |

|      |  |    |
|------|--|----|
| 7.3  | Work Breakdown Structure .....   | 27 |
| 7.4  | Project Management Control System .....  | 27 |
| 7.5  | Project Bootstrapping.....   | 28 |
| 8.   | Financial and Business Controls.....   | 28 |
| 8.1  | Financial and Business Controls.....   | 28 |
| 8.2  | Financial and Progress Reporting .....   | 28 |
| 8.3  | Status Reports .....   | 28 |
| 8.4  | Subcontracting Controls.....   | 29 |
| 9.   | Performance Assessment and Quality Assurance.....  | 29 |
| 9.1  | Quality Assurance Plan.....  | 29 |
| 9.2  | Performance Assessment Plan .....  | 30 |
| 10.  | Network Plan .....   | 31 |
| 10.1 | Network Description.....   | 31 |
| 10.2 | Services Provided by FutureGrid Network.....   | 32 |
| 10.3 | Service Levels.....  | 33 |
| 11.  | Software Plan .....  | 35 |
| 11.1 | Information Services.....  | 35 |
| 11.2 | Experiment-Management Services .....   | 35 |
| 11.3 | Actuating Services: Test-Bed Management, Experiment Initiation, Collection of<br>Experimental Data, Storage of Data for Later Retrieval and Use..... | 36 |
| 12.  | Systems Integration and Transition to Operational Status Plan.....   | 37 |
| 12.1 | PY1 – Basic Implementation .....   | 37 |
| 12.2 | PY2 and Beyond – Operational Phase .....   | 38 |
| 13.  | Risk Management Plan .....   | 39 |
| 13.1 | Management Risks.....  | 39 |
| 13.2 | Operations and Facilities Risks.....   | 40 |
| 13.3 | Software Risks .....   | 41 |
| 14.  | Interface Agreements .....   | 42 |
| 14.1 | FutureGrid Access Mechanisms .....   | 42 |
| 14.2 | Networking .....   | 43 |
| 14.3 | Security .....   | 43 |
| 14.4 | Accounts .....   | 43 |
| 14.5 | Services .....   | 44 |

|  |    |
|--|----|
| 14.6 Operations .....  | 46 |
| 14.7 Support.....  | 46 |
| 14.8 Management.....   | 47 |
| 15. Cybersecurity Plan .....                                     | 47 |
| 15.1 System and Infrastructure Security .....                    | 48 |
| 15.2 Software Security .....                                     | 50 |
| 15.3 Incident Response .....                                     | 50 |
| Appendix A: FutureGrid Project Plan Milestone Schedule .....     | 52 |
| Appendix B: WBS Dictionary .....                                 | 59 |
| Appendix C: FutureGrid Project Schedule .....                    | 69 |
| Appendix D: Projected Annual Cost by WBS .....                   | 83 |
| Appendix D.1: Projected Annual Cost by WBS NSF-IU Breakdown..... | 84 |
| Appendix E. FutureGrid <b>DRAFT</b> Staffing Plan.....           | 85 |
| Appendix F. IU Notice to Subcontract Recipients.....             | 87 |

## 1. Executive Summary

This Program Execution Plan describes the implementation of *FutureGrid—an experimental, high-performance grid test-bed*. This Plan discusses the science supported by FutureGrid and how we will attract and assist researchers; the organizational roles of the individual and institutional participants; the management plan; project deliverables; our plans for engaging Minority Serving Institutions; project management, budget, and reporting processes; performance assessment; networking, software, and systems; risk management; interface agreements; and cybersecurity. In the appendices, we discuss project milestones and scheduling, along with details of the work breakdown structure. The FutureGrid project will create deliverables in six categories: a facility, software, educational materials, scientific data and knowledge, better-educated students, and careful reporting and dissemination of its accomplishments.

Clouds are challenging assumptions about grid computing. The goal of FutureGrid is to support the research that will invent the future of distributed, grid, and cloud computing. FutureGrid will be a cyberinfrastructure for the development of new approaches to scientific applications and for distributed computing research. It will be operated as a single unified instrument, and is perhaps more unlike the existing TeraGrid resources than any other resource funded through the Track II program.

For computer and computational science researchers developing middleware – grid software, cloud software, and new types of middleware yet to be named – FutureGrid will provide a rich and flexible test-bed. FutureGrid will enable rigorous scientific experiments in grid and cloud computing, resulting in significant extensions to existing software as well as new software development. While FutureGrid is a test-bed environment, it will be crucial that the FutureGrid network perform as expected. The FutureGrid network will follow standard best practices for maintenance and operations to ensure high availability and predictability for the resource.

Quality assurance is an integral part of the FutureGrid project. Our goal is responsiveness to user requirements and the evolving collaborative development and delivery of the environment that supports the testing and evaluation needs of FutureGrid users. Our Performance Assessment Plan will consist both of continual feedback on the quality of services and of more formal quarterly and annual reporting and review processes.

PI Geoffrey Fox will lead overall management of the project. Fox, Executive Director Craig Stewart, and co-PIs Kate Keahey, Warren Smith, Jose Fortes, and Andrew Grimshaw will form the FutureGrid executive committee. Two advisory committees and six operational committees will provide additional input into the management of FutureGrid. The funded university participants (Indiana University, University of California – San Diego, University of Chicago, University of Florida, University of Southern California, University of Tennessee – Knoxville, University of Texas at Austin – Texas Advanced Computer Center, and University of Virginia) will each have representatives on the relevant committees. Unfunded university participants (Purdue University and Technische Universitaet Dresden) and a private sector partner (GWT-TUD GmbH) will also play key roles in FutureGrid.

## 2. Science Plan

### 2.1 Motivation and Purpose

Innovation and discovery in science and engineering have been revolutionized by the ever-growing confluence of application science, computational science, and informatics. Increasingly sophisticated large-scale simulations and rapidly growing data sets have led to concepts such as eResearch, eScience, and Cyberinfrastructure. These advances are built on distributed computing, parallel computing, and their integration. With rapidly expanding network, storage, and computing requirements, new application science systems will require the development of new and more innovative enabling cyberinfrastructure.

As science-driven needs are growing, we face a crucial time in academic distributed computing. Academia is lagging behind industry in distributed and parallel computing research as Google, Microsoft, and others invest billions of dollars in technology and infrastructure. Clouds are challenging assumptions about grid computing. Multicore computing means ubiquitous parallel computing. Researchers increasingly require advanced computational science applications for use in nontraditional fields and by nontraditional groups.

This Program Execution Plan describes the implementation of *FutureGrid—an experimental, high-performance grid test-bed*. The goal of FutureGrid is to support the research that will invent the future of distributed, grid, and cloud computing. FutureGrid will support the development and early use in science of new technologies at all levels of the software stack: from networking to middleware to scientific applications. This test-bed will enable dramatic advances in science and engineering through collaborative evolution of science applications and related software. Table 1 outlines many of the general types of grid and computational science experiments that we plan to support via FutureGrid.

The computational science community has a strong need for facilities that enable a more scientific approach to comparison and evaluation of distributed computing software. The critical element of the science plan for FutureGrid is that it will enable rigorous, repeatable experiments in middleware and distributed computing, facilitating the sort of exactitude for distributed computing systems and performance analysis that has long characterized parallel performance analysis. Repeatability is based on the ability to instantiate a particular environment, in isolation from outside interference, with a particular and repeatable set of initial conditions. Networks will generally be dedicated to particular experiments, and, when network impairments are involved in an experiment, they will be generated through use of a network impairment device, allowing for repeatability. Data stored for any given experiment will include the system images in which an experiment was performed, along with the software actually used and input data. Hence, FutureGrid will be a cyberinfrastructure for the development of new approaches to scientific applications and for distributed computing research.

We expect that the activities that will take place within FutureGrid will be primarily experiment-based, driven by an experiment plan or involving steps that may be viewed as an experiment plan. That plan may be very basic: instantiate a particular environment and let a researcher debug an application interactively, or very sophisticated: instantiate a particular environment and run a pre-specified set of tasks. A direct outcome of this experiment-centric approach is that it will lead to a collection of software images and experimental data that will prove a tremendous resource for application and computational sciences.

| Use case   | Required to fulfill use case   |
|--|--|
| Testing a new networking protocol or topology, application layer overlays, and peer-to-peer networks         | Ability to build system images and propagate them through a test environment<br>Dedicated time in an isolated test environment, with prescribed and repeatable levels of load and error conditions   |
| HCI researchers testing end-to-end productivity of grid computing systems                                    | Variety of software and hardware environments allowing presentation of multiple systems and user interfaces  |
| Testing grid or cloud, particularly end-user applications  | Specify a grid or cloud environment and run applications in that environment; compare with other environments<br>Prepare applications for deployment on commercial systems (cloud or grid)<br>Test a complex workflow, which requires a heterogeneous hardware mix |
| Creating a cloud front end linked to a grid and its resources to enable scientific applications and gateways | Cloud test environment, ability to link to one or potentially many different hardware architectures as back end  |
| Developing data-intensive applications   | Link data sources to a grid environment specified by the developer, possibly including supported workflow tools—for example LIGO data flow, medical images, or sensor data   |
| Testing optimization of different layers of parallelism via grid, cloud, and many-core programming models    | Grid or cloud test environment that includes systems representing varying levels of core counts per processor  |
| Comparing grid middleware implementations and standards compliance   | Persistent endpoints for grid interoperability testing<br>Test-bed to compare grid operating environments  |
| Testing new authentication or authorization mechanism  | Ability to run a persistent authentication server in test environment or link to one at the researcher's lab   |
| Hardening of middleware or science application   | Security vulnerability ("simulated attack") test service<br>Simulated job load<br>For network- or grid-centric applications, ability to simulate latency, inject errors into network, etc.   |
| Testing performance of applications on non-x86-64 architectures  | For resource providers, the ability to place non-x86-64 architectures in a multiuser environment<br>For application developers, the ability to test applications on non-x86-64 architectures to evaluate code performance  |

**Table 1. Experimental grid test-bed requirements matrix. Common needs across all of these use cases include the ability to (1) specify a test environment in advance and use it during a scheduled period of time and (2) create an appropriate record of an experiment, save it securely, and retrieve it reliably in the future.**

For computer and computational science researchers developing middleware – grid software, cloud software, and new types yet to be named – FutureGrid will provide a rich and flexible test-bed, and will be a platform for computer and computational scientists to use for developing new network, distributed, grid, and cloud applications; and for rigorously evaluating new approaches at all levels, from application science down through the layers of technology to networking.

We will support application science directly and indirectly. Application scientists and software developers can develop and prove new approaches to delivery of their applications. Such applications can then be migrated to other production cyberinfrastructure facilities, enabling better support and delivery of end-user science capabilities to the U.S. research community. We will support network, grid, cloud, and distributed computing directly by providing an environment that supports computer and systems research that will lead to improved cyberinfrastructure that indirectly supports application science. Dedicated networking and 24 x 7 monitoring will provide a secure environment in which new applications can be safely developed, tested, and hardened.

## 2.2 *Early Science Experiences*

FutureGrid is perhaps more unlike the existing TeraGrid resources than any other resource funded through the Track II program. The TeraGrid has added new large systems, experimental hardware, and high-throughput systems. However, no experimental test-bed system has ever been part of the TeraGrid (except when the TeraGrid experimental itself). The history of the TeraGrid suggests that it can take considerable time for the U.S. research community to recognize and make good use of a novel type of resource within the TeraGrid. With a team that brings together some of the very best of leaders in academic grid and cloud research, it will be tremendously important to achieve a good balance between ensuring that FutureGrid is well used early on, and having so much of FutureGrid's use come from our own team that we create a perception that FutureGrid serves the FutureGrid team first and foremost. Table 2 lists researchers who have already indicated that they would like to use the FutureGrid test-bed during its initial year of testing and development and their science projects. We have identified primary points of contact within the FutureGrid team for these projects and are committed to working closely with these researchers to ensure they are up and running as soon as possible.

We are at an unusual time in the evolution of cyberinfrastructure in the United States. There is tremendous excitement about the potential of cloud computing. There are new and interesting developments in grid middleware. There are two planning grants funded to define the future of TeraGrid eXtreme Digital (TeraGrid XD). In this environment, it will be particularly important to generate some early and interesting, high-impact, useful results that leverage FutureGrid as a grid and cloud test-bed and could not have been achieved without a dedicated test-bed. It is hard in advance to predict which particular projects listed in Table 2, or other new projects, will lead to the most exciting results early on. However, many of the projects listed have the potential to be very interesting and have high impact quickly, particularly the projects that involve running Science Gateway front-ends in cloud environments, work on grid interoperability and standards compliance, and undergraduate education activities from U. Florida. The next sections describe three potential use cases that may prove particularly exciting during the early stages.

### 2.2.1 Teaching Grid Software

The future of scientific computing is distributed computing. With the emergence of dense multicore and similar architectures for personal computing, the proliferation of smart devices and sensors on the real-time Internet, and the evolution of large-scale production instruments, it is important to provide a new and forward-looking teaching environment that integrates seamlessly with large-scale cyberinfrastructure. Achieving this goal requires programmers and domain scientists who understand grid, distributed, and parallel programming. Current production cyberinfrastructure such as the TeraGrid is not ideal for teaching – a student might even crash a grid while learning to program it. In addition, students learning to program grids may introduce real and severe security vulnerabilities; even seconds of exposure may be all it takes for a malicious actor to gain unauthorized access to a computing system. In order to let students program in a safe and encapsulated environment, we will create an environment that will allow the creation of a virtual grids in which students can experience the full complexity of grid computing for writing and debugging grid software, allowing students to use a variety of cloud and grid computing environments. FutureGrid is designed to support precisely this sort of virtual grid creation for educational purposes, by creating a grid sandbox for student use. We will also provide a variety of pre-configured grid and cloud environments. Students will be excited to do real grid programming on real grids, a capability currently unavailable.



| Researcher                                 | Institution                        | Planned use   |
|--|------------------------------------|---|
| <b>Researchers outside FutureGrid team</b> |                                    |   |
| Bruce Berriman                             | California Institute of Technology | Astronomy, applications such as Montage   |
| Minaxi Gupta                               | IU                                 | Reliability and security in peer-to-peer networks   |
| Shantenu Jha                               | LSU                                | SAGA (genome analysis) implementation with MapReduce  |
| Mark Miller                                | SDSC                               | Cloud-based version of CIPRES phylogenetics portal  |
| Randy Katz                                 | UC Berkeley                        | Development of multisite cloud management tools   |
| Sriram Krishnan                            | UCSD                               | NEESit (earthquake engineering)   |
| Renato Figueiredo                          | U. Florida                         | Cloud environments as educational collaborative tools—education in virtual machines, virtual networks   |
| Klaus Schulten                             | U. Illinois Urbana Champaign       | NAMD and VMD  |
| Ann Duin                                   | U. Minnesota                       | Graduate student classes in operating systems   |
| Kelvin Droegeheimer                        | U. Oklahoma                        | Testing of LEAD on clouds   |
| Ben Berman and Peter Laird                 | USC                                | Epigenomic applications   |
| James Bower                                |                                    | K–12 interactive learning in parallel and distributed computing   |
| <b>Researchers within FutureGrid team</b>  |                                    |   |
| Geoffrey Fox                               | IU                                 | Comparative performance analysis of Hadoop, Eucalyptus, Azure   |
| Craig Stewart and Rich Knepper             | IU                                 | Comparative performance of Windows HPC Server vs. Linux on NSF benchmarks   |
| Marlon Pierce                              | IU                                 | U.S. Geological Survey application delivery from Eucalyptus   |
| Song/Smith/Campbell                        | Purdue                             | Determining effective strategies for building, monitoring, and maintaining multi-institutional grids  |
| Song/Smith/Campbell                        | Purdue                             | Identifying and improving mechanisms for implementing campus grids to maximize technology investments (particularly instructional labs)           |
| Song/Smith/Campbell                        | Purdue                             | Investigate the interface between cloud technologies and Condor VM  |
| Song/Smith/Campbell                        | Purdue                             | Explore mechanisms that promote grid interoperability (e.g. TeraGrid and OSG)   |
| Chris Jordan                               | TACC                               | Distributed/WAN file system testing   |
| Warren Smith                               | TACC                               | Evaluate metascheduling tools and algorithms  |
| Edward Walker                              | TACC                               | Testing adaptive load balancing algorithms in MyCluster   |
| Matthias Mueller                           | TUD                                | Performance analysis of VMs and applications with new versions of Vampir  |
| Kate Keahey                                | UC                                 | New applications in distributed Nimbus cloud  |
| Shava Smallen                              | UCSD                               | Integration and scalability testing for Inca grid monitoring software   |
| Shava Smallen                              | UCSD                               | Testing for UC (University of California) grid portal software  |
| Jose Fortes                                | UFL                                | Management of BLAST and BLAST-related jobs on single and multiple clouds  |
| Jose Fortes                                | UFL                                | Development of multicloud networking management tools   |
| Ewa Deelman                                | USC                                | Explore data management issues within scientific workflows on the cloud   |
| Ewa Deelman                                | USC                                | Quantify overheads in scientific workflows running on the cloud   |
| Ewa Deelman                                | USC                                | Compare execution of workflows running on the TeraGrid versus the cloud   |
| Andrew Grimshaw                            | UVA                                | Deploy Genesis II infrastructure and perform usability studies comparing middleware exposed through the file system with “traditional” middleware |
| Andrew Grimshaw                            | UVA                                | Test Genesis II middleware on large sets of workflows, where each workflow instance consists of both sequential and low-degree parallel phases    |
| Andrew Grimshaw                            | UVA                                | Perform interoperability tests between existing standards-compliant middleware stacks such as GridSAM, Unicore, g-lite, and Genesis II            |

**Table 2. Planned research activities for PY1.**

### 2.2.2 Maximizing Computing Power for Public Safety Applications

In our global society, we are confronted with global public health and public safety applications that require the integration of compute power on a global scale. In many cases, the ability to apply computing power in real time to ad-hoc computational problems is often a matter of life and death. For example, Alex Vespignani's GIEaM (Global Epidemic and Mobility model) application has been used throughout the year to predict the spread of H1N1. GIEaM is one example of a kind of application that can benefit from as much computing resource as can possibly be made available in urgent situations. One way to enable access to real-time computing power needed for such a demanding application is to integrate different grids. Unfortunately, instead of a single huge grid, we currently have many large grids. We will test a deployment of GIEaM across U.S. and European test grids and effectively increase its responsiveness to global health issues. Such a demonstration can be done via FutureGrid and the French Grid5000 experimental grid. Within the United States, FutureGrid is the only national facility that can provide a persistent connection point to European grids. Once testing is complete, it will be possible to run GIEaM across the U.S. TeraGrid and large European grids such as EGEE (Enabling Grids for E-science). In this way, we can prepare for a day when the health of millions of people might be improved by running simulations across the TeraGrid and European grids simultaneously.

### 2.2.3 Dynamic Allocation of Computational Resources for Genomic Research

The TeraGrid's current configuration is optimized for scalable parallel computation done largely in a batch mode. Increasingly, however, many scientists are finding that they need cyberinfrastructure to handle the deluge of data from scientific instruments such as genome sequencers. We will use the new virtualization capabilities of FutureGrid to implement the CIPRES genomic analysis systems in a cloud front end with dynamically allocated computational resources as a back end. CIPRES is the software being used to develop a new tree of life. The aim of this project is literally to compare the genome sequence of every species against every other species. The integration of clouds and computational systems will provide a new and remarkable way of using computational resources different from the traditional batch-job work style. FutureGrid in particular supports open-source cloud environments that are compatible with the Amazon cloud environments, meaning that application front ends developed with FutureGrid can be used on other installations of these open-source tools or on Amazon's commercial cloud resources. The variety of computational resources available within FutureGrid makes it possible for scientists to test scalable supercomputers, clouds, and other specialized resources as back ends. As a result, we can create portable, next-generation tools for supporting life science research that will aid CIPRES and also provide an example architecture that can be used by other life scientists and data-centric sciences in general.

## 2.3 *Attracting and Selecting Interesting and Valuable Research*

As mentioned above, it will be very important to get off to a quick start, both because there is much valuable and time-sensitive science to be supported and because the greatest benefit to the U.S. research community will be delivered by providing early access to users. The issue of attracting scientists to use FutureGrid can be broken down into three categories – attracting computational science researchers; attracting domain scientists; and attracting educators to use FutureGrid in instructional settings.

### 2.3.1 Attracting Computational Scientists

Computational scientists will be attracted to FutureGrid through a variety of mechanisms – talks and posters at conferences, articles in such publications as HPCwire, announcements on NSF, TeraGrid, and Open Science Grid web pages, etc. We believe it will be relatively straightforward to generate interest among computational scientists especially right now, with so many claims and counterclaims regarding performance, efficacy, and ease of use of many new grid and cloud environments. We believe there is tremendous interest in being able to do grid and cloud research with the sort of rigor that in past years has characterized parallel scalability research. We believe that the computational science community will be highly motivated to conduct high-quality research in a configurable test-bed. We plan to enhance motivation to use FutureGrid by making it convenient for researchers to deposit results in open repositories such as <http://www.myexperiment.org/>. We plan to demonstrate FutureGrid basic functionality at SC09, which should generate considerable interest very quickly. We also intend to invite CLUE awardees to experiment with FutureGrid and, via the appropriate program officer (Jim French), invite scientists who made unfunded proposals to CLUE to try FutureGrid as well. Also, according to minutes of past National Science Board meetings, there were a total of three proposals submitted in response to solicitation NSF 08-573 proposing an experimental grid test-bed. Once a public announcement has been made, IU will ask that NSF staff route invitations to use FutureGrid to the leadership of the other two proposing teams, in hope and anticipation that participants in those teams will be interested in using FutureGrid.

### 2.3.2 Attracting Domain Scientists

We expect it to be somewhat more challenging to attract domain scientists to FutureGrid, and a rational expectation is that domain scientists will be most interested in testing applications running in small-scale cloud environments. We have already implemented a few domain science applications, including a U.S. Geological Services GIS application – certainly not a typical TeraGrid application. We plan to disseminate information through domain science conferences and workshops. For example, Beth Plale has already invited Fox and Stewart to give a very brief presentation at a September 2009 workshop on “Cloud Computing and Collaborative Technologies in the Geosciences.” We will do so provided a public announcement of FutureGrid makes this appropriate. We believe two factors will be critical in attracting domain scientists to FutureGrid: making the process of applying for usage simple and sending domain scientists to domain science conferences to discuss the value of the facility to the science domain. We plan to do both.

### 2.3.3 Attracting Educators

We believe that the key to attracting educators will be having early exemplars of successful use of FutureGrid in education, and high-quality curriculum materials that educators can adapt and reuse. We expect that the U. Florida group will develop and enhance curriculum materials that can be distributed and reused by other educators. Similarly, U. Minnesota has indicated strong interest in being an early adopter. We expect to develop materials at IU and share them with U. Minnesota and with IU’s MSI partner institutions (as described in Section 6).

## 2.4 *Detailing Capabilities of FutureGrid*

The current methods of displaying TeraGrid resources within the TeraGrid user portal are highly effective for production use of TeraGrid resources. However, we believe this format is not likely

to be optimal for presenting the capabilities of FutureGrid. We plan to display the capabilities of FutureGrid through a combination of maps and tables showing the full extent of the planned system, accessible from a portal specific to FutureGrid. We will document the development and implementation of FutureGrid resources by “graying out” resources planned but not yet available. We will similarly disseminate the software capabilities available at each node of FutureGrid by listing available and planned (grayed out) capabilities. This should make it straightforward for researchers to both understand what capabilities are available and plan future experiments by understanding what capabilities are planned but not yet deployed.

## 2.5 *Identifying Research Problems and Allocating Usage*

The TeraGrid allocation process represents the outcome of 20 years of experience and refinement. However, while it is regarded as much improved, it is still perceived as difficult to negotiate by many. Rather than start with a complicated acceptance process, we plan during the first half of PY1 to ask researchers to write a resource request in the form of an explanation of the experiment they wish to perform with FutureGrid and a list of the resources and software capabilities they will need. We anticipate that this resource request may initially be as short as two pages. We will also ask that requestors attach a standard NSF-format two-page biosketch for the PI and any co-PIs. This approach should minimize barriers to adoption while at the same time allowing us to learn over time how best to structure later, more formal, resource requests.

Co-PI Andrew Grimshaw (who chairs the FutureGrid Science Advisory Board as described in section 4.3.1) and a committee he will advise will review these initial resource requests. Throughout the award, the FutureGrid team will reserve 10% of the FutureGrid resource for use by the FutureGrid team, to be allocated as PI-discretionary time. To ensure good records of activities, all projects using time under the PI’s 10% discretionary time will submit a resource request in the same format as any other user.

We will generally be heavily biased in favor of fulfilling early requests in particular, in the belief that by so doing we can best facilitate the development of new computational tools (middleware and application software), and best learn how to develop more refined and formal templates for resource requests during the latter half (PY 3 and 4) of the project. The later evolution of the allocation process is described in detail in section 12.2.

## 3. **Organizational Roles**

Organizational roles are described below for each institution with a focus on the categories of hardware systems management; software development, networking, and security; performance analysis; user support; education and outreach; and project management.

### 3.1 *Funded University Participants*

**Indiana University (IU).** IU will be responsible for the overall management of the FutureGrid project. As the home institution of the PI, IU is ultimately responsible for the success of FutureGrid. The largest suite of hardware within FutureGrid will be located at Indiana University, and IU will chair the hardware management team. IU will also lead the interactions between FutureGrid (as an instrument within the TeraGrid) and the TeraGrid (and in the future TeraGrid XD) as a whole. Particular areas of responsibility include

- **Systems administration and network management.** IU will lead the hardware management of FutureGrid. In particular, the chair of the Systems Administration and Network Management Committee will be located at IU (the inaugural chair will be David Hancock). IU will host an IBM iDataPlex, a Cray system, a shared memory system to be identified, and a small Cell BE cluster. IU will also host a centrally located Spirent network impairment device.
- **Software adaptation, implementation, hardening, and maintenance.** IU will chair the Software Adaptation, Implementation, Hardening, and Maintenance Committee (the inaugural chair will be Gregor von Laszewski). IU will lead in software development, particularly as regards development of initial tools for instantiating environments on request. IU will lead the creation of the FutureGrid user portal.
- **Performance analysis.** IU will, via matching funds, manage a subcontract with GWT-TUD GmbH for support of Vampir for users of FutureGrid.
- **Support.** IU will provide operational coordination for user support. This will include provision of information to users via an online Knowledge Base, 24 x 7 telephone support (emergency only outside 8am to 8pm Eastern Time), a trouble-ticket management system for FutureGrid, and operational activities between FutureGrid and the TeraGrid (and later TeraGrid XD) as a whole.
- **Training, education, and outreach services (TEOS).** IU will provide operational coordination for training, education, and outreach services, and will develop and deliver training materials and tutorials.
- **Project management.** PI Fox will lead this project overall. Executive Investigator Stewart will also serve in a leadership role. IU will be responsible for overall project management, including management of any and all reporting required by the NSF or TeraGrid (and later TeraGrid XD) leadership. An IU staff member, initially Gary Miksik, will be devoted 0.5 FTE to project management of FutureGrid (Appendix E). IU will also be the primary U.S. liaison to the German D-Grid project (<http://www.d-grid.de>).

**University of California – San Diego (UCSD).** UCSD will lead the Performance Analysis Committee, participate in performance analysis activities, adapt and deploy software for systems monitoring software to aid the operation of FutureGrid, and host an IBM iDataPlex system that will be part of FutureGrid. Particular areas of responsibility include

- **Systems administration and network management.** UCSD will operate an IBM iDataPlex system as part of FutureGrid.
- **Software adaptation, implementation, hardening, and maintenance.** UCSD will adapt and extend Inca as part of the FutureGrid management software.
- **Performance analysis.** UCSD will chair the Performance Analysis Committee (inaugural chair will be Shava Smullen).
- **Support.** UCSD will provide advanced support for Inca, and second-tier support for users of the hardware resource located at UCSD. UCSD will also prepare Knowledge Base entries relevant to Inca.
- **Training, education, and outreach services.** UCSD will provide training materials relevant to use of Inca and the performance analysis tests developed by UCSD and used within FutureGrid.

- **Project management.** UCSD will participate in project management and reporting so as to ensure that reports are submitted on time and requests for information from the NSF or advisory boards are fulfilled.

**University of Chicago (UC).** UC will be responsible for support of Nimbus for FutureGrid users, will host an IBM cluster as part of FutureGrid, and will participate in TEOS activities. Particular areas of responsibility include

- **Systems administration and network management.** The University of Chicago will host and manage an IBM iDataPlex as part of the FutureGrid test-bed environment.
- **Software adaptation, implementation, hardening, and maintenance.** The University of Chicago will be responsible for deployment of Nimbus within FutureGrid.
- **Support.** The University of Chicago will provide advanced support for Nimbus and second-tier support for users of the hardware resource located at UC. UC will also prepare Knowledge Base entries relevant to Nimbus.
- **Training, education, and outreach services.** UC will provide training materials relevant to use of Nimbus within FutureGrid.
- **Project management.** UC will participate in project management and reporting so as to ensure that reports are submitted on time and requests for information from the NSF or advisory boards are fulfilled. UC will also serve as FutureGrid's liaison to the European Grid5000 project. As one of the co-PIs, Kate Keahey will participate in the leadership of the FutureGrid project.

**University of Florida (UF).** UF will be responsible for deployment of ViNe (Virtual Network) within FutureGrid, particularly its use to support educational and training activities. Particular areas of responsibility include

- **Systems administration and network management.** UF will manage a small cluster primarily for purposes of developing and testing ViNe.
- **Software adaptation, implementation, hardening, and maintenance.** UF will enhance the current integration of ViNe, integrating the routing layer with Nimbus so that it is easy to create self-configuring virtual networks and virtual appliances within Nimbus, and then expanding the capabilities of ViNe to also function within Eucalyptus.
- **Support.** UF will provide advanced support for ViNe. UF will also prepare Knowledge Base entries relevant to ViNe.
- **Training, education, and outreach services.** UF will apply virtual-appliance- and social-networking-based systems developed at UF (and based on ViNe) to facilitate dissemination of FutureGrid software for education, development, and testing. In particular, UF will develop self-learning educational modules that will allow teachers and students to download grid software within a virtual appliance and experiment with it on small-scale local hardware. UF will develop a how-to tutorial and support a social networking group related to FutureGrid on Facebook.
- **Project management.** UF will participate in project management and reporting so as to ensure that reports are submitted on time and requests for information from the NSF or advisory boards are fulfilled.

**University of Southern California (USC).** USC will support use of Pegasus within FutureGrid, and work with other developers of FutureGrid software to implement experiments within FutureGrid as workflows executed via Pegasus. Particular areas of responsibility include

- **Systems administration and network management.** USC has no responsibilities in this area.
- **Software adaptation, implementation, hardening, and maintenance.** USC will support use of Pegasus by FutureGrid users. USC will integrate Pegasus and other experiment-management systems so that grid experiments can be implemented as a workflow within Pegasus.
- **Performance analysis.** Pegasus will be used to collect and consolidate data resulting from performance analysis experiments, and USC will provide second-tier support for researchers who want to do performance experiments with Pegasus particularly.
- **Support.** USC will provide advanced support for Pegasus. USC will also prepare Knowledge Base entries relevant to Pegasus.
- **Training, education, and outreach services.** USC will participate in outreach activities. These activities will take two forms. First, because Pegasus is capable of integrating and automating complicated workflows, it has considerable potential applicability to a broad array of domain sciences that may or may not currently be heavy users of the TeraGrid. A key component of USC's outreach will encourage domain scientists who are not currently users of the TeraGrid to experiment with Pegasus, creating workflows that automate work now done by hand. In addition, as a leading woman computer scientist, Ewa Deelman will be involved in activities that focus on encouraging women to pursue careers in computing and science, technology, engineering, and mathematics (STEM) disciplines.
- **Project management.** USC will participate in project management and reporting so as to ensure that reports are submitted on time and requests for information from the NSF or advisory boards are fulfilled.

**University of Tennessee – Knoxville (UTK).** UTK will develop and support tools for benchmarking grid applications. Particular areas of responsibility include

- **Systems administration and network management.** No responsibilities.
- **Software adaptation, implementation, hardening, and maintenance.** No responsibilities except those specifically related to performance analysis.
- **Performance analysis.** UTK will support PAPI on FutureGrid systems. UTK will modify the existing HPC Challenge benchmark test for execution across FutureGrid (and other grid computing environments). Furthermore, UTK will develop a new test-bed suite specifically designed for grid test-beds called the Grid Benchmark Challenge, based on the general model of HPCC.
- **Support.** UTK will develop Knowledge Base entries related to PAPI, HPCC in grid environments, and the Grid Benchmark Challenge. UTK will provide second-tier support for FutureGrid users making use of these tools.
- **Training, education, and outreach services.** UTK will develop training materials relevant to PAPI, HPCC for grid environments, and the Grid Benchmark Challenge.
- **Project management.** UTK will participate in project management and reporting so as to ensure that reports are submitted on time and requests for information from the NSF or advisory boards are fulfilled.

**University of Texas at Austin – Texas Advanced Computer Center (TACC).** TACC will host a Dell blade cluster as part of the dedicated FutureGrid hardware environment, and provide

access to other systems located at TACC as appropriate. TACC will participate in the development of the FutureGrid user portal, and lead development of test harness software. Particular areas of responsibility include

- **Systems administration and network management.** TACC will manage a Dell blade cluster as part of the hardware dedicated to FutureGrid. In addition, as appropriate and as allocated by the TeraGrid Resource Allocation Committee (TRAC), TACC will make its Ranger and Spur systems available as part of grid experiments. This is not expected to include on-the-fly rebuilding of either Ranger or Spur. However, either or both systems might be used in an experiment using experimental grid workflow systems. For example, a grid experiment might involve computing at scale with Ranger as one element of a larger test. Or, a workflow system test might involve visualization with Spur as one element of a workflow.
- **Software adaptation, implementation, hardening, and maintenance.** TACC will participate in the development of the FutureGrid user portal. TACC will also be responsible for the creation and support of a test harness for executing experiments on FutureGrid.
- **Performance analysis.** No specific responsibilities other than development of the test harness to be used in performance analysis experiments.
- **Support.** TACC will develop Knowledge Base entries related to the test harness, and provide second-tier support for FutureGrid users making use of the test harness.
- **Training, education, and outreach services.** TACC will develop class materials that involve use of FutureGrid.
- **Project management.** TACC will participate in project management and reporting so as to ensure that reports are submitted on time and requests for information from the NSF or advisory boards are fulfilled. As one of the co-PIs, Warren Smith will participate in leadership of FutureGrid.

**University of Virginia (UV).** UV will support use of Genesis II, Unicore, and EGEE software on FutureGrid. UV will also serve as the primary FutureGrid liaison to the Open Grid Forum and grid-standard working groups. Particular areas of responsibility include

- **Systems administration and network management.** No responsibilities.
- **Software adaptation, implementation, hardening, and maintenance.** UV will support deployment of Genesis II, Unicore, and EGEE software on the dynamically configurable FutureGrid nodes. In addition, UV will maintain stable and ongoing installations of Genesis II on a small number of Intel-based nodes at IU for interoperability testing.
- **Performance analysis.** No responsibilities beyond second-tier support of software mentioned above.
- **Support.** UV will develop Knowledge Base entries related to Genesis II, Unicore, and EGEE software, and provide second-tier support for FutureGrid users making use of these tools.
- **Training, education, and outreach services.** UV is already developing educational materials regarding Genesis II, and these will be useful to users of FutureGrid.
- **Project management.** UV will participate in project management and reporting so as to ensure that reports are submitted on time and requests for information from the NSF or advisory boards are fulfilled. As one of the co-PIs, Andrew Grimshaw will participate in



leadership of FutureGrid, particularly by chairing the Science Advisory Board and the User Advisory Board.

### *3.2 Unfunded University Participants*

**Purdue University (PU).** PU will provide a 96-node high-throughput cluster for use within FutureGrid, provide access to a small cluster of FPGAs, and serve as a backup site for hosting hardware. Particular areas of responsibility include

- **Systems administration and network management.** Purdue University will provide a 96-node high-throughput cluster as part of the FutureGrid test-bed connected to FutureGrid systems via the I-light network. Purdue University will also provide access to a small cluster of FPGAs or other specialized-processor systems.
- **Software adaptation, implementation, hardening, and maintenance.** Purdue University will support use of Condor and BOINC on the high-throughput cluster.
- **Project management.** Purdue will participate in project management and reporting so as to ensure that reports are submitted on time and requests for information from the NSF or advisory boards are fulfilled.

**Technische Universitaet Dresden (TU-D).** TU-D will provide limited use of one of its high performance computing systems for transatlantic grid testing, will participate in performance analysis activities, and will serve as a liaison to the German D-Grid project. Particular areas of responsibility include

- **Systems administration and network management.** TU-D will provide limited access to its SGI Altix system (neptun) or its successors, for transatlantic grid testing activities.
- **Software adaptation, implementation, hardening, and maintenance.** None other than those related to performance analysis.
- **Performance analysis.** TU-D will participate in analysis of network and grid performance between the United States and Germany, and collaborate with FutureGrid in trying to establish a suite of official SPEC benchmark applications. TU-D will also provide early access to Vampir and VampirTrace software that will particularly support performance analysis within virtual machines (VMs).
- **Project management.** TU-D will participate in project management and reporting so as to ensure that reports are submitted on time and requests for information from the NSF or advisory boards are fulfilled. TU-D will serve as the primary point of contact with the German D-Grid project.

### *3.3 Private Sector Partners*

**GWT-TUD GmbH.** GWT-TUD GmbH will, under a contract with Indiana University funded as part of its match commitment, provide support for FutureGrid users making use of Vampir and VampirTrace software during PY2–4.

## **4. Management Plan**

### *4.1 Overall Management Structure*

Fox will lead overall management of the project. Fox with the co-PIs will form the FutureGrid executive committee. Stewart will serve as executive director for the project and oversee

operations. Gregor von Laszewski will serve as Software Architect and oversee all technical aspects of software development and integration. Together they will be responsible for integration of activities across committees in their respective areas (operations and software). FutureGrid will be operated as a single unified instrument. We will most particularly not be replicating the current TeraGrid Forum model in which participating sites are semi-autonomous.

Operational management will operate via a group of key management personnel and a suite of committees, each charged with leading a particular area of FutureGrid activities. Decisions made by these committees will be binding upon FutureGrid as a whole. In general, each committee will have participants from relevant participating institutions (e.g., all institutions hosting hardware as part of FutureGrid will participate in the Systems Administration and Network Management Committee, but those not hosting hardware generally will not). Figure 1 shows the organizational structure of FutureGrid and the various committees that will lead management of FutureGrid.

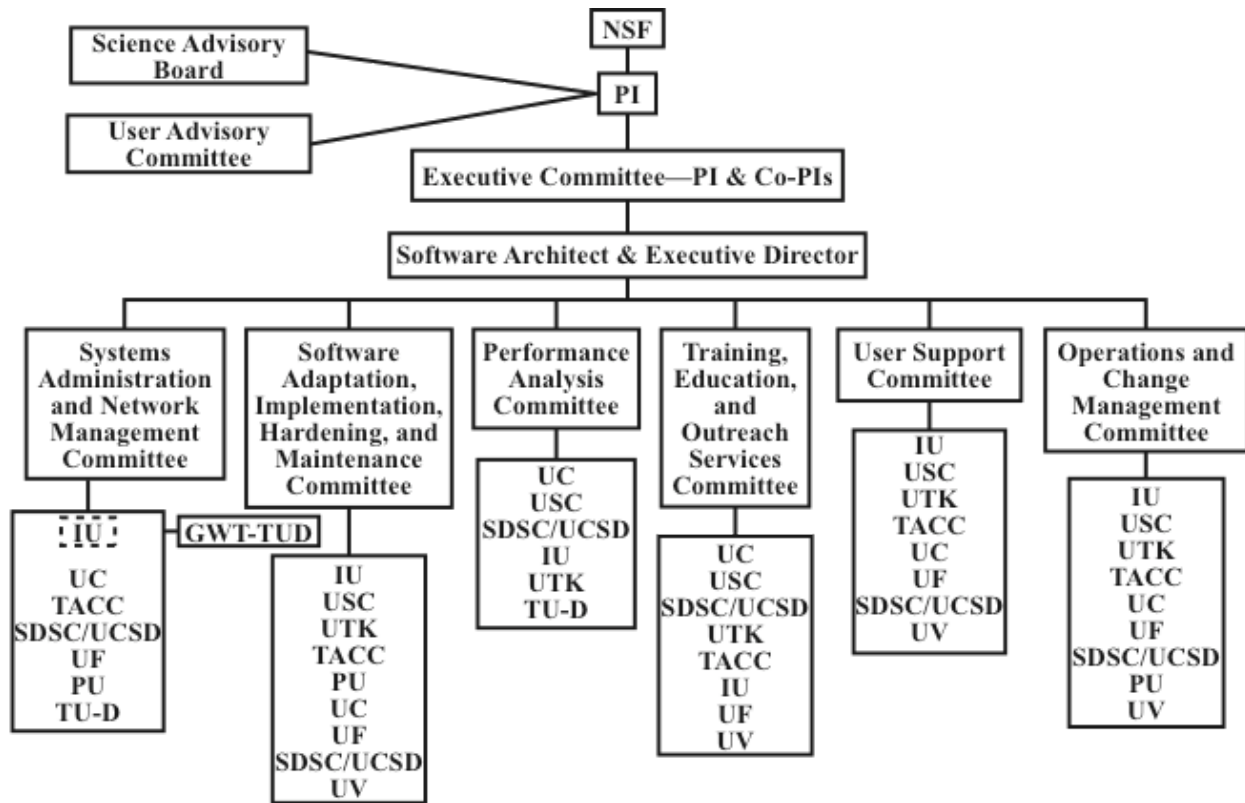


Figure 1. FutureGrid organizational structure.

#### 4.2 Key Management Personnel

**PI.** Geoffrey Fox will be the PI, and have overall responsibility for the project as a whole. Fox will be the final arbiter of any decisions that cannot be reached by a consensus approach.

**Executive Director.** Craig Stewart will serve as executive director, responsible particularly for operational management of FutureGrid.

**Co-PIs.** Kate Keahey, Warren Smith, Jose Fortes, and Andrew Grimshaw will serve as co-PIs; each has a particular leadership role within FutureGrid.

**Chief Architect.** Gregor von Laszewski (who joined IU on 22 July 2009) will serve as the chief architect for FutureGrid.

**Site Lead.** Rich Knepper will serve as site lead for FutureGrid, and will be responsible for interfacing between FutureGrid and the TeraGrid (and later TeraGrid XD).

**Project Manager.** Gary Miksik will serve 0.5 FTE as project manager for FutureGrid, and have management of the WBS, preparation of reports, and collection of responses to requests for information from the NSF as his primary job responsibilities.

### *4.3 Committee Structures*

Each committee will be expected to have one high-quality, well-organized meeting per month. Additionally, the Operations and Change Management Committee will have weekly status check/change management meetings.

#### 4.3.1 Advisory Committees

**Science Advisory Board.** A Science Advisory Board (SAB) with strong links to science users and TeraGrid XD will meet at least twice per year. One meeting will be in January to finalize annual work plans. The other meeting will be held in early summer as a mid-year progress review. This committee will consist primarily of grid computing experts and will be chaired and convened by co-PI Andrew Grimshaw.

**User Advisory Committee.** An advisory committee of users will be convened distinct from the SAB, to provide direct feedback from FutureGrid users and nonusers to FutureGrid leadership. The inclusion of grid researchers who are not users of FutureGrid is considered particularly important; they will advise FutureGrid regarding the needs we are not meeting successfully. This committee will be chaired and convened by co-PI Andrew Grimshaw.

#### 4.3.2 Operational Management Committees

**Systems Administration and Network Management Committee.** This committee will be responsible for all matters related to systems administration, network management, and security. David Hancock of IU will be the inaugural chair of this committee.

**Software Adaptation, Implementation, Hardening, and Maintenance Committee.** This committee will be responsible for all aspects of software creation and management. The FutureGrid software architect will chair this committee.

**Performance Analysis Committee.** This committee will be responsible for coordination of performance analysis activities. Shava Smallen of UCSD will be the inaugural chair of this committee.

**Training, Education, and Outreach Services Committee.** This committee will coordinate Training, Education, and Outreach Service activities and will be chaired by co-PI Jose Fortes.

**User Support Committee.** This committee will coordinate the management of online help information, telephone support, and advanced user support. Jonathan Bolte of IU will chair this committee.

**Operations and Change Management Committee.** This committee will be responsible for operational management of FutureGrid, and is the one committee that will always include at least one member from every participating institution, including those participating without funding.

This committee will be responsible for tracking progress against the work breakdown structure (WBS), preparing reports, managing finances, and general coordination. This committee will also include the chairs of every other committee within FutureGrid. This committee will also serve as a Change Control Board (CCB), meeting biweekly to review and approve changes before they are implemented. (The CCB will be available to meet more often to handle ad hoc requests.) FutureGrid Project Manager Gary Miksik will chair this committee. This committee will also oversee use of the discretionary 10% of FutureGrid resource usage reserved for the FutureGrid team.

**Executive Committee.** This committee is the second highest authority within the FutureGrid management structure, second only to the PI himself.

#### *4.4 Consensus Management Process*

Committees will operate according to a consensus process. Rather than having “yea/nay” votes, there will be four votes: Strongly in favor; in favor; opposed; strongly opposed. Consensus is declared when there is a plurality of votes in the combined categories of “strongly in favor” and “in favor” and there are no “strongly opposed” votes. This process generally works well when there is an across-the-board commitment to success and spirit of collaboration, as we expect within FutureGrid committees. When it is impossible to reach consensus, committee chairs will render final decisions. Conflicts may be escalated to the executive committee. Consensus may be reached there, and when consensus even there is impossible, the PI will render a final decision. As a general rule, we expect decisions to be made quickly and do not expect and will not tolerate stalemates in discussion.

#### *4.5 Maintaining, Refreshing, and Executing the Project Vision*

The proposal to create FutureGrid and this Project Execution Plan set out a vision for a cyberinfrastructure for distributed, grid, and cloud computing research. It will be important to maintain that vision and as appropriate refresh it. The Operations and Change Management committee will include representatives of all participating institutions. It is this group that will be most responsible, on a day in–day out basis, for ensuring that project execution is consistent with project vision. This Project Execution Plan will serve as an initial statement of the vision for FutureGrid. The vision will be updated and refreshed annually by the SAB and executive committee. Input for such updating will come from meetings of the Science Advisory Board, the User Advisory Board, BOFs at TeraGrid and Supercomputing conferences, and discussions with the TeraGrid Science Advisory Board (and its successors), as well as NSF staff. We plan a formal update of the PEP every year.

## **5. Project Deliverables**

The FutureGrid project will create deliverables in six categories: a facility; software; educational materials; scientific data and knowledge; better educated students; and careful reporting and dissemination of its accomplishments.

### *5.1 Facility*

FutureGrid will be an unparalleled national-scale grid and cloud test-bed facility that includes a total of at least nine computational resources – six of which are new – from at least three vendors (IBM, Cray, Dell, and one to be determined), four different types of file systems, and a network

that can be dedicated to perform repeatable experiments in isolation, including a network impairment device for repeatable experiments under a variety of predetermined network conditions (see Table 3). Also, FutureGrid will be connected to an archival storage system.

| System type                                 | # CPUs      | # Cores     | TFLOPS    | RAM (GB)     | Secondary storage (TB) | Default local file system | Site |
|---|-------------|-------------|-----------|--------------|------------------------|---------------------------|------|
| <b>Dynamically configurable systems</b>     |             |             |           |              |                        |                           |      |
| IBM iDataPlex                               | 256         | 1024        | 11        | 3072         | 335*                   | Lustre                    | IU   |
| Dell PowerEdge                              | 192         | 1152        | 12        | 1152         | 15                     | NFS                       | TACC |
| IBM iDataPlex                               | 168         | 672         | 7         | 2016         | 120                    | GPFS                      | UC   |
| IBM iDataPlex                               | 168         | 672         | 7         | 2688         | 72                     | Lustre/PVFS               | UCSD |
| <b>Subtotal</b>                             | <b>784</b>  | <b>3520</b> | <b>37</b> | <b>8928</b>  | <b>542</b>             |                           |      |
| <b>Systems not dynamically configurable</b> |             |             |           |              |                        |                           |      |
| Cray XT5m                                   | 168         | 672         | 6         | 1344         | 335*                   | Lustre                    | IU   |
| Shared memory system TBD                    | 40**        | 480**       | 4**       | 640**        | 335*                   | Lustre                    | IU   |
| Cell BE Cluster                             | 4           |             |           |              |                        |                           |      |
| IBM iDataPlex                               | 64          | 256         | 2         | 768          | 5                      | NFS                       | UF   |
| High Throughput Cluster                     | 192         | 384         | 4         | 192          |                        |                           | PU   |
| <b>Subtotal</b>                             | <b>552</b>  | <b>2080</b> | <b>21</b> | <b>3328</b>  | <b>10</b>              |                           |      |
| <b>Total</b>                                | <b>1336</b> | <b>5600</b> | <b>58</b> | <b>10560</b> | <b>552</b>             |                           |      |

Table 3. High-level hardware specifications for systems to be included as part of the FutureGrid test-bed and available no later than 1 October 2010. \*Indicates shared file system. \*\*Best current estimate

Figure 2 shows a schematic map of FutureGrid.

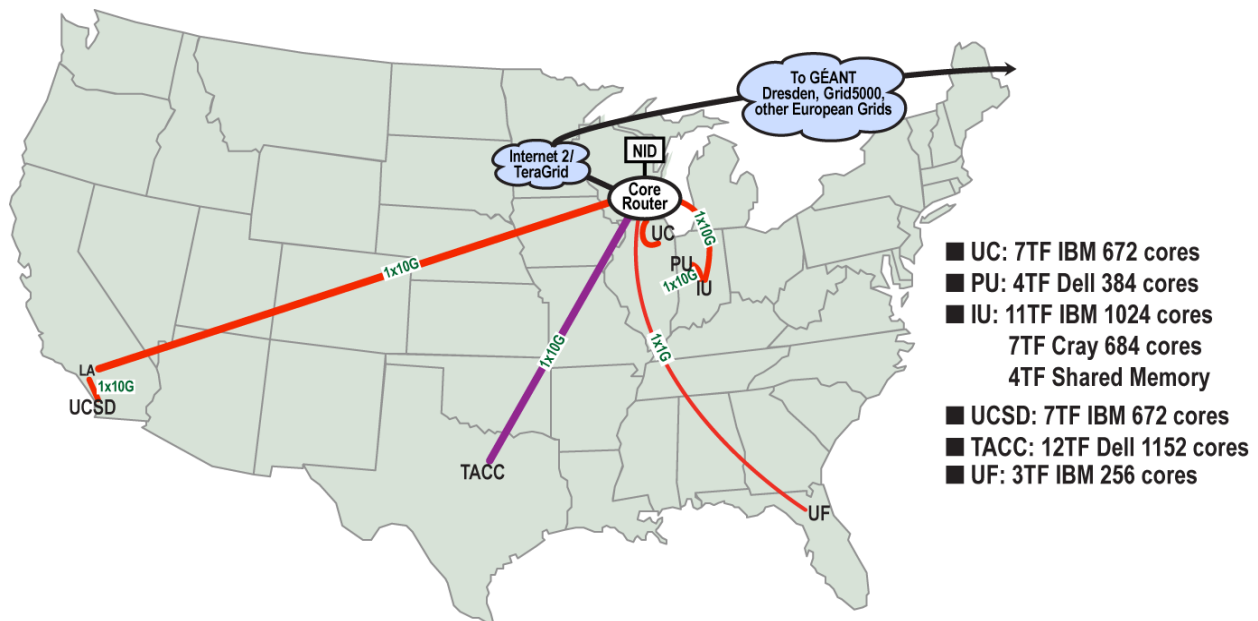


Figure 2. Schematic map of FutureGrid. All network links will be dedicated at initiation of operations, except the link to TACC. That link will initially be shared; we expect to be able to implement dedicated links as needed dynamically once TeraGrid Phase III (XD) is implemented. (NID = Network Impairment Device)

## 5.2 Software

As a result of the FutureGrid project, there will be significant extensions to existing software as well as new software developed. We will extend existing open-source software to create an

open-source, integrated suite of software to instantiate and execute grid and cloud experiments. This software suite will instantiate a prespecified software environment within a computing grid, perform an experiment, and collect the results. The software suite will include tools for instantiating a test environment, which we anticipate will include Torque, MOAB, xCAT, bcfg, and Pegasus. Pegasus' workflow engine supports dynamic environments. Inca will be extended to support cloud test-bed environments. We will develop a portal, in keeping with TeraGrid standards and practices, for preparing such experiments.

We will develop a grid version of the widely used HPCC benchmark suite, and then develop a new Grid Benchmark Challenge application suite.

### *5.3 Educational Materials*

We will develop and openly disseminate curricular materials that will encourage and enable use of FutureGrid. Such materials will also be useful as a basis for developing other new curricular materials in grid and cloud computing. [Note: section 6 details several educational activities particular to engagement of minority serving institutions (MSIs).]

### *5.4 Scientific Data and Knowledge*

FutureGrid will enable rigorous scientific experiments in grid and cloud computing. We will store output of these experiments in an archival storage system at IU. The FutureGrid team will enable and encourage researchers who use FutureGrid to store experimental results in open, public repositories.

FutureGrid will aid international understanding of grid and cloud computing by partnering with other experimental grids such as Grid5000. FutureGrid will thus contribute to (as opposed to create) new standards for international grid interoperability.

In addition to the specific scientific data and knowledge that is created, FutureGrid will nurture a culture of rigor in grid and cloud computing comparable to the traditions of scientific approaches to scalable computing.

### *5.5 Better Educated Students*

Through the net effects of its educational activities and the provision of research platforms used by graduate students, FutureGrid will better educate a cohort of students in computational sciences. Through outreach efforts and targeted recruitment, FutureGrid should in particular enable students from traditionally underserved groups to pursue careers in computing and STEM disciplines.

### *5.6 Reports, Presentations, and Published Works*

FutureGrid will produce a clear record of its activities and outcomes. Reporting to the NSF will create an objective, thorough record of the accomplishments of FutureGrid. Presentation materials will be created and widely disseminated (primarily in forms that may be reused and repurposed). Additionally, FutureGrid participants and users will create a body of research published through peer-reviewed scientific journals and as technical reports.

## 6. Minority Serving Institution Engagement Plan

The FutureGrid team will leverage extensive, pre-existing resources and infrastructure at Indiana University to involve MSIs in our project. This allows us to offer virtual clusters and test-beds focused on teaching and developing FutureGrid applications. FutureGrid will provide the following resources to MSIs. These capabilities are a consequence of expected operations and require no additional effort. In order to make MSIs aware of FutureGrid capabilities available to them, we will engage in an outreach program (Section 6.4) for which we are well prepared (Section 6.2). Activities will include providing resources for MSI faculty to teach systems programming on individual machines and clusters as well as preconfigured, dynamically instantiated environments for teaching parallel programming, web programming, grid and cloud programming, and computational science (Section 6.1).

Principal Investigator Geoffrey Fox has an established track record of working with MSIs. Similarly, Dr. Jose Fortes and his colleague Dr. Renato Figueiredo have expertise in use of social networking techniques for engaging individuals from traditionally underserved groups. We note the distinction between engaging with MSIs as opposed to engaging with a few students from MSIs. This plan is for engagement at the institutional level. It is based on two key strategies – leveraging the MSI Cyberinfrastructure Empowerment Coalition (MSI-CIEC) and using social networking tools.

### 6.1 Goals of Engagement

FutureGrid is designed to support fundamental research in parallel and distributed computing. We will build FutureGrid on virtualization and cloud computing technologies. In the process of operating FutureGrid, we will acquire a library of virtual machines encapsulating many important distributed computing research efforts: Condor, Globus, Apache Hadoop, OpenMPI, and Genesis II, to name a few. These libraries of virtual machines and virtual clusters will provide an easily installed and evaluated platform for classroom and other educational uses. The core software products underlying FutureGrid (Eucalyptus, Nimbus, Pegasus, and others) also represent important distributed computing research efforts. We will build upon the extensive MSI outreach resources at Indiana University and the virtual library of FutureGrid to provide instructional resources for MSI faculty to teach modern distributed computing.

Our goals for collaborating with MSIs are the following:

- Teaching faculty how to use FutureGrid resources (virtual machines and virtual clusters) to teach basic distributed computing, systems programming, and system administration in the classroom. FutureGrid provides a secure sandbox that will allow each student to have his/her own test-bed in isolation from other students and operational facilities.
- Providing MSI faculty with preconfigured environments for teaching parallel, web, distributed, and grid computing.
- Enabling teaching and research collaborations between MSI institutions and experts in grid and cloud technologies and research.
- Teaching faculty how to build test-bed versions of FutureGrid out of resources at their institutions for classroom use.
- Teaching students how to use FutureGrid tools through internships.
- Ultimately, ensuring that computational sciences in particular and STEM disciplines in general have the benefit of the talents of the best and brightest individuals. Conversely,

we wish to engage such students through FutureGrid and expose them to a scientific instrument shaping the future.

- ~~Ensure that every post-secondary student in the US has educational opportunities that enable them to pursue a technology-centric career of their choice.~~

Through our established connections with MSIs and established outreach programs, FutureGrid is ideally positioned to support larger national activities that seek to ensure that U.S. students of all racial, ethnic, and economic backgrounds wishing to pursue technical or scientific careers will have access to resources and educational material. As we move into an operational phase with FutureGrid, we will create memoranda of understanding (MOUs) with MSIs regarding specifics of partnership activities to be undertaken relative to FutureGrid.

## 6.2 *Leverage of MSI Cyberinfrastructure Empowerment Coalition (MSI-CIEC)*

Fox is currently a principal member of the MSI Cyberinfrastructure Empowerment Coalition (MSI-CIEC), which has been funded by the NSF CI-TEAM and other awards. MSI-CIEC's primary theme is to "teach the teachers" at MSIs so that they can incorporate cyberinfrastructure into their research and involve students and staff at their home institutions. MSI-CIEC's current principal activity is the organization of Cyberinfrastructure Days at various MSIs. These daylong workshops feature prominent speakers who discuss the application of cyberinfrastructure to research and education.

In addition to the MSI-CIEC, Fox and the FutureGrid team will work closely with Maureen Biggers, Indiana University's assistant dean for diversity and education. Biggers' qualifications include acting as project manager for the National Science Foundation's Broadening Participation in Computing Alliance for the Advancement of African-American Researchers in Computing, and as a member of the leadership team for the National Center for Women and Information Technology. We will work with Biggers to organize outreach and pursue REU funding to bring MSI students to IU for summer internships and to coordinate education and training workshops.

Finally, FutureGrid will involve students from Historically Black Colleges and Universities (HBCUs) through Indiana University's STEM Initiative (<http://www.stem.indiana.edu/>). This program provides travel, housing, and support for HBCU students to intern at Indiana University during the summer. We particularly expect to engage the MSIs listed in Table 4, with which Indiana University has already established formal collaborative agreements.

| <b>Institution</b>                | <b>Location</b> |
|-----------------------------------|-----------------|
| Alabama A&M                       | Normal, AL      |
| Bennett College for Women         | Greensboro, NC  |
| Clark Atlanta University          | Atlanta, GA     |
| Hampton University                | Hampton, VA     |
| Jackson State University          | Jackson, MS     |
| Langston University               | Langston, OK    |
| Morgan State University           | Baltimore, MD   |
| Morehouse College                 | Atlanta, GA     |
| Xavier University                 | New Orleans, LA |
| Tennessee State University        | Nashville, TN   |
| North Carolina Central University | Durham, NC      |
| Clark Atlanta University          | Atlanta, GA     |



**Table 4. Minority Serving Institutions with which Indiana University has a formal collaborative agreement, and which we expect to engage in using FutureGrid.**

### 6.3 *Leveraging Social Networking Technologies*

U. Florida will apply virtual appliance and social networking systems developed at U. Florida (<http://www.grid-appliance.org>, <http://www.socialvpn.org>) to facilitate the dissemination of the grid test-bed software for education, training, and development. This will allow MSI educators and students to quickly (within minutes) gain hands-on access to a system that has the same software stack of the grid test-bed but runs on their own resources – without worrying about software installation, configuration, or the time taken to request and process an account.

The system we propose allows an individual or groups of users to easily deploy an ad-hoc virtual private network of virtual machines that would run the same software that runs in the grid test-bed. All they need to do is download a VM image that runs out-of-the-box in a free VM monitor (e.g., VM Player or VirtualBox), create a group in a social network infrastructure (e.g., Facebook), and turn on the appliances to create an ad-hoc virtual cluster. This enables interesting usage scenarios in education and training.

### 6.4 *Workshops and Tutorials Specific to MSI Engagement*

We will support our engagement goals through a series of workshops and tutorials. These will be offered both online and face to face. Online material will include both live and archived material. Topics to be covered are discussed in Section 6.1.

Specific deliverables:

- We will offer self-guided tutorials on an ongoing basis, with the first tutorial available by the end of the first year, and at least one tutorial per year in Program Years 2–4.
- We will offer at least one face-to-face, daylong tutorial onsite at MSIs. These will be based on our “CI Days” workshop series from MSI-CIEC. All material will be archived and made available through the FutureGrid web site and related resources. These will be offered via established connections with MSI institutions in Indiana, Kentucky, Illinois, Ohio, and Michigan.
- Some FutureGrid resources past their lifecycle will be donated to MSIs for teaching.

## 7. Project Budget and Work Breakdown Structure

This section includes a summary of Level 1 WBS Definitions, description of the methodology and assumptions used for estimating budget components, and description of the project management control system.

Additional details are provided in several appendices, as follows.

- Appendix A: FutureGrid Project Plan Milestone Schedule
- Appendix B: WBS Dictionary
- Appendix C: Project Schedule
- Appendix D: Projected Annual Cost by WBS

Some of the particularly important aspects of the FutureGrid plan are detailed in sections 10 through 12.

## 7.1 Project Budget

The anticipated start date of the project is 1 October 2009. The projected annual costs by cost type are shown in Table 5. The budget distributed across the WBS is shown in Appendix D. The budgets by year are estimates and, while they may change from year to year, the total cost to the NSF is fixed at \$10,100,000. Payment terms for hardware and network invoices are detailed in the respective vendor contracts. These contracts will be signed by the authorized purchasing agent and will remain on file during the project period for this grant. The budget completes on 30 September 2013 (end of project), unless the Cooperative Service Agreement is modified.

| NSF Funding by Category             | Cost (\$M)    | FY 2010       | FY 2011       | FY 2012       | FY 2013       |
|-------------------------------------|---------------|---------------|---------------|---------------|---------------|
| Salaries and fringe benefits        | \$4.88        | \$1.21        | \$1.21        | \$1.22        | \$1.23        |
| Hardware                            | 1.83          | 1.83          | —             | —             | —             |
| Networking                          | 0.13          | 0.11          | 0.01          | 0.01          | 0.01          |
| Travel                              | 0.14          | 0.03          | 0.04          | 0.03          | 0.03          |
| Other                               | 0.35          | 0.08          | 0.09          | 0.09          | 0.07          |
| Indirect costs                      | 2.67          | 0.72          | 0.65          | 0.64          | 0.65          |
| NSF funding                         | 10.1          | 4.10          | 2.00          | 2.00          | 2.00          |
| IU cost share                       | 5.7           | 2.32          | 0.97          | 1.40          | 0.98          |
| <b>Grand total NSF + cost share</b> | <b>\$15.8</b> | <b>\$6.33</b> | <b>\$2.98</b> | <b>\$3.41</b> | <b>\$3.00</b> |

Table 5. Projected annual costs by cost type in millions of dollars.

## 7.2 Methodology and Assumptions Used for Estimating Budget Components

Staff salaries were based on the salary of existing staff expected to take a role in the FutureGrid project. In some cases, we have a catch-22 situation: There are staff who must be in an appointed position to perform the work anticipated, but permission to hire people and place them in appointed positions will not be given by budget offices until an award instrument has been finalized with the NSF. In these cases, salaries are based on expected salaries (and in some select cases salary expectations that have been prospectively discussed with existing, unappointed staff). Almost all of the institutions participating in FutureGrid have previous experience with TeraGrid or PACI programs. The amount of time (both elapsed calendar time and person-months) required to perform tasks is based primarily on the prior experience of participating institutions with the TeraGrid. U. Florida is the only participating institution without extensive experience with the TeraGrid. Their work is largely an extension of funded work they have already been doing, or unfunded work done in collaboration with UC. Their time estimates are thus based on that prior work.

Fringe benefits and indirect costs were calculated using the published and approved institutional fringe benefit and indirect cost rates at each participating institution. Indirect costs were based on the current Indirect Cost rate negotiated between each participating institution and the NSF.

The basic hardware and network costs were derived first from some basic beliefs about what was needed for an effective grid test-bed: It should include a diversity of hardware types, reflecting both systems currently represented in the TeraGrid and systems not yet so represented. We also believed that it was essential to have a test-bed that was really distributed, as opposed to some form of simulation. We additionally believe that it was important that the aggregate number of processor cores be at least 5000 (the initial core target for the French Grid5000 project). Final hardware and network cost figures were taken from actual vendor quotes.

Based on the capabilities we thought essential for an effective grid test-bed, and the characteristics described above, we came to the conclusion that a grid test-bed including the

capabilities we believed important could not be done within the \$10,000,000 budget set out in the NSF solicitation. As a result, IU has contributed substantial institutional match to bring the total committed budget for the project to \$15,810,702. In the end, some labor is likely underestimated, as seems typical for TeraGrid-related activities. Some participating institutions expect to contribute more effort than indicated in formal documents, but this is not indicated formally due to difficulties in negotiating match agreements with local research offices. In addition, travel is likely underestimated, again in keeping with TeraGrid traditions. The increase in travel suggested by the NSF largely remedied this, along with revisions to the Virginia subcontract. Some participating institutions may need to supplement their travel budgets; all understand this.

### *7.3 Work Breakdown Structure*

FutureGrid tasking is broken into six (6) major categories, which constitute the Level 1 WBS, defined below.

- **(WBS 1.0) Hardware.** This category encompasses all activities related to the procurement, installation, use, and management of computer resources at each FutureGrid site, including contractual acceptance benchmarks. Each hardware vendor's tasking will be tracked at a separate Level 2 WBS.
- **(WBS 2.0) Networks.** This category encompasses all activities related to network connectivity between all sites, including the procurement, installation, and ongoing support of network devices. Each network vendor's tasking will be tracked at a separate Level 2 WBS.
- **(WBS 3.0) Software.** This category encompasses all activities related to the design, development, and deployment of the various software modules and components to be made available for use in FutureGrid. Each specific software component is tracked at a separate Level 2 WBS.
- **(WBS 4.0) Support.** This category encompasses those activities directly related to the ongoing support provided to FutureGrid users, including help desk, knowledge base, and advanced consulting services.
- **(WBS 5.0) Training, Education, and Outreach (TEO).** This category encompasses the activities related to how FutureGrid information gets disseminated to both its users and the general population. Specific TEO tasks are tracked at separate Level 2 WBS.
- **(WBS 6.0) Project Management (PM).** This category targets all activities related to the planning, management, and coordination of the other project elements to assure the NSF investment will be successful. Specific PM tasks are tracked at separate Level 2 WBS.

### *7.4 Project Management Control System*

The project management control system will be based on Microsoft Project software. The project manager will implement a project management control system to manage FutureGrid's project scope, cost, and schedule. Using Microsoft Project, the project manager will manage the work breakdown structure (WBS) and alert the executive committee of any cost or schedule variances. The project manager will also be responsible for tracking the status of all deliverables and being aware of any slipping deliverables so the executive committee can be alerted and resources can be reallocated as necessary. In cases where decisions are needed more urgently, the PI will make the decision and inform the executive committee of the issue via e-mail or telephone.

## 7.5 *Project Bootstrapping*

The FutureGrid Project Plan Milestone Schedule is presented in Appendix A and assumes a start date of 1 October 2009. We plan an aggressive schedule for hardware installation and implementation of dedicated network links. Early in Program Year 1, we will issue an open call for early science users, and we will begin the development process with a series of research input activities (surveys, calls, meetings, etc.). We expect the annual SC conference to be a highlight of demonstration activity each November. Each December, we will survey researchers who are using FutureGrid to determine what aspects of the overall service are working well, what needs improvement, and what new features are needed. Each January, we will plan activities from January to September of the following calendar year.

## 8. Financial and Business Controls

### 8.1 *Financial and Business Controls*

All financial and business controls and standards in place at Indiana University will be followed. Internal audit and internal management oversight will be used to monitor the project as required. Formal oversight of all cooperative service agreements is the responsibility of the IU Office of Research Administration (ORA). The IU Accounts Payable department will make payments to partner organizations from approved invoices.

### 8.2 *Financial and Progress Reporting*

Project tasking will be supported by a combination of Excel worksheets and Microsoft Project plans. Budgets and actual costs will be collected from official financial accounts established in the IU Financial Information System (FIS) and will reflect the project's overall WBS in reports to the NSF. The chief operating officer (COO) of the Pervasive Technology Institute at Indiana University will oversee the execution of all project budgets. The FutureGrid project manager is responsible for reporting project progress and project financials to the NSF and for ensuring that invoices submitted for payment by partner organizations are correct.

### 8.3 *Status Reports*

The FutureGrid project manager is responsible for working with the FutureGrid team in preparation of quarterly and annual reports. Quarterly and annual status reports will be approved by the principal investigator and submitted to the NSF Program Office for final approval.

#### 8.3.1 Annual Reports

- Give the key accomplishments in the prior year.
- Provide a comprehensive financial report.
- State project changes that occurred during the year, including but not limited to schedule variance, cost variance, schedule adjustments, prioritization of the next year's tasks, and management of reserve allocations.
- Include the annual metrics report.
- Summarize major risk handling activities accomplished in the prior year and describe the current risk status.
- Fulfill all information requests made by the NSF or advisory committees.

### 8.3.2 Quarterly Reports

The project manager will also submit quarterly status reports to the NSF Program Office. These reports document major accomplishments and project changes and include a financial report. The quarterly reports will fulfill information requests made by the NSF or advisory committees.

## 8.4 *Subcontracting Controls*

The FutureGrid project manager, with expertise provided by the Indiana University Purchasing Department, is responsible for planning, executing, and tracking all procurements required for the completion of the project. The FutureGrid PI and project manager are responsible for general coordination with all of the organizations involved in procurement planning to ensure that (1) procurement requirements are properly defined; (2) major procurements are included in the project schedule to identify required delivery dates and to allow for adequate lead time for all phases of the project; and (3) procurements are budgeted properly such that the project baseline is consistent with the procurement requirements and schedule.

FutureGrid partner organizations are classified as *subrecipients* at Indiana University. The Office of the Vice President for Research Administration at Indiana University has issued an *Important Notice* documenting all subrecipient processes and the responsibilities for subrecipient monitoring. A copy of this notice is in Appendix F and will be the basis for all partner organization interaction on the FutureGrid project.

## 9. Performance Assessment and Quality Assurance

### 9.1 *Quality Assurance Plan*

Quality assurance is integral to the FutureGrid project. Beyond the responsibility to deliver stable and consistent availability of what will be quite a complex and heterogeneous hardware and software environment, FutureGrid has a unique larger role in supporting the quality of the TeraGrid as a whole, in that it provides a test environment for the software stack and middleware that underlie current and emerging grid and cloud environments.

FutureGrid will consist of distributed and varied hardware systems. The initial quality control for each of those systems will be part of the agreements covering the acceptance of the hardware from the various vendors. Ongoing evaluation of the hardware will be supported through the use of monitoring software including Inca. Irregularities observed from these monitoring probes will be entered into our incident response system and escalated as required through resolution.

The implementation of the middleware software that will provide, for example, the dynamic configuration and virtualization functions will be supported in an agile and iterative fashion. There will be continual feedback in the form of incident reporting and ongoing user needs elicitation. Our goal is responsiveness to user requirements and the evolving collaborative development and delivery of an environment that supports the testing and evaluation needs of FutureGrid users.

The FutureGrid team expects to work in close collaboration with TeraGrid XD's Technology Audit and Insertion Services (TAIS) group upon its establishment. This should be a synergistic relationship in which TAIS will establish methodologies through which FutureGrid will monitor the quality of its services, and FutureGrid will be able to allocate resources to TAIS users to test

and evaluate the quality of TeraGrid software. These combined efforts have a common goal of increased value and improved user experience.

## 9.2 *Performance Assessment Plan*

The FutureGrid Performance Assessment Plan (PAP) will consist of both continual feedback on the quality of services and more formal quarterly and annual reporting and review processes. The goals of the PAP will be the measurement of usage and performance metrics in absolute terms against set goals and milestones, trending of these metrics and user satisfaction over time, and ongoing improvement of services based on feedback.

Performance metrics will include those specified by the NSF. We expect also to develop measures of utilization that are relevant specifically to the test-bed, including the time required to initialize and begin experiments. Trouble-ticket and event-tracking systems will be monitored on an ongoing basis for trends in volume and time to resolution. We expect to meet or exceed TeraGrid standards for problem resolution.

Performance against milestones related to project build out will be reviewed regularly and an escalation process will be established to bring slippage to the attention of project and TeraGrid leadership for resolution. Each quarter, we will tally achievement of milestones, categorizing deliverables at each level of the WBS into one of three categories: achieved, less than one quarter late, and more than one quarter late. For milestones less than one quarter late, we will ask the person responsible for accomplishment of that milestone whether the milestone is reachable within the next quarter without reconsideration of project plans. We will reconsider any milestone that is projected to be more than one quarter late and prepare a remediation and risk mitigation plan. Any milestone that is and remains more than one quarter late will be subject to reconsideration and there will be an update to a remediation and risk mitigation plan for each such milestone at the end of each quarter in which the milestone remains incomplete.

Feedback on quality of services and service needs will be sought through regular surveys to assess user satisfaction. Surveys will also be offered to non-TeraGrid users engaged in grid computing research to solicit a broader spectrum of input for service improvement. An anonymous mechanism will be available for any feedback and process improvement suggestions. Surveys and interviews will be completed with approval of Indiana University's Internal Review Board, in compliance with IU Human Subjects Research policy, so that we may publish or otherwise disseminate all results of such surveys and interviews.

Beginning on 1 October 2011, FutureGrid will participate in all formal TeraGrid reporting and review processes for TeraGrid and XD, assuming a continuation of the current combined project-wide reporting and review process. FutureGrid will employ processes established by TeraGrid XD to measure science impact. Publications enabled by FutureGrid allocations will be tracked and science highlights and images related to FutureGrid research will be made available regularly.

As part of our quarterly and annual reporting, we will report work products produced according to their NSF categorizations (web pages, peer-reviewed technical articles, etc.) We will focus our attention most heavily on tracking original technical papers in peer-reviewed primary journals.

All collected data will be used by FutureGrid leadership to assess and evaluate processes and services on an ongoing basis and as part of a more formal annual service-evaluation process. FutureGrid's oversight committee will review performance data and make recommendations to

FutureGrid. Minor process improvements will be implemented on an ongoing basis based on impact as anticipated by FutureGrid leadership. Major changes in FutureGrid services will be coordinated through the TeraGrid forum (through the end of the current TeraGrid awards) and then through the TeraGrid CMS management organizations.

## **10. Network Plan**

The FutureGrid network will provide for interconnections among FutureGrid participants and access to the FutureGrid network impairments device. Figure 3 shows the FutureGrid network topology.

### *10.1 Network Description*

#### 10.1.1 Core and National Backbone

The core of the FutureGrid network will consist of a FutureGrid network core router, located at the Starlight facility in Chicago. This will be a Cisco 6500 series router. A series of dedicated links will connect the FutureGrid core router with the FutureGrid participants at UF, IU, and UCSD. TACC will use a shared access via TeraGrid for initial connectivity to FutureGrid.

The FutureGrid network will use a 10-Gigabit Ethernet dedicated lambda from NLR to connect UCSD to the core router. This will take a path from the 710 N Lakeshore Drive facility to the NLR location in Los Angeles. This will cost \$68,716 per year and will have a 4-year term. It can be renewed for an additional year at the same annual rate.

The FutureGrid network will use three 10-Gigabit Ethernet ports from NLR on the NLR Los Angeles FrameNet switch to allow multiple separate access to the dedicated lambda from Los Angeles to the 710 N Lakeshore Drive facility. This will give full access to the 10-Gigabit Ethernet available, and this will not be shared with any other FrameNet traffic. This will cost a one-time fee of \$25,920 for three ports.

A 1 Gigabit dedicated bandwidth VLAN from NLR will be used to connect UF to the core router. This VLAN is provided by National Lambda Rail FrameNet network, from the FrameNet node in the 111 N Canal Street, Chicago, location to the FrameNet node in the NLR location in Jacksonville, Florida at a cost of \$17,520 recurring annually with a four-year term. This can be renewed for an additional year at the same annual rate.

An NLR lambda will be used to connect the FutureGrid core router at 710 N Lakeshore Drive to the NLR FrameNet node at 111 N Canal Street at a cost of \$17,179 recurring annually with a four-year term. This can be renewed for an additional year at the same annual rate.

The FutureGrid network will make use of 1 10-Gigabit Ethernet port on the NLR FrameNet switch in the 111 North Canal Street location in Chicago. This will initially be used to provide the 1 Gigabit dedicated bandwidth VLAN for connectivity to the University of Florida. This will also allow other potential FutureGrid users with NLR FrameNet connections to connect additional VLANs to the FutureGrid core router.

#### 10.1.2 Site Networking

For UCSD, CENIC will provide a 10-Gigabit Ethernet dedicated lambda from the UCSD system to the NLR location in Los Angeles.

For UF, FLRnet will provide a 1-Gigabit Ethernet dedicated VLAN from the UF system to the NLR location in Jacksonville. This will be provided at a capital equipment expense of \$5,523 and no annual recurring charges.

For IU, the dedicated 10-Gigabit Ethernet network connection to Chicago will be contributed as match to the NSF at a value of \$54,780 in Program Year 1 capital expense with no annual recurring charges. Purdue will also connect through this same dedicated connection by leveraging the IP-Grid network to Indianapolis.

For UC, Starlight will provide a 10-Gigabit Ethernet dedicated lambda from the UC system to the location of the core router in downtown Chicago. This will be provided at an expense of \$30,000 in Program Year 1 and an annual recurring fee of \$6,000 with a four-year term. This can be renewed for an additional year at the same annual rate.

For TACC, their existing 10-Gigabit Ethernet connection to Chicago will be utilized at no additional cost to the NSF. The link from TACC will not be dedicated, sharing their existing connection to the TeraGrid or using TACC's redundant TeraGrid connection when a dedicated link is required by an experiment.

#### 10.1.3 Network Impairments

A Spirent H10 XGEM network impairment simulator will be collocated with the FutureGrid core router to simulate the types of network impairments that might be encountered on a production network. This device was chosen because it is the only device on the market today that can provide full network impairment simulation of 10Gbps flows of any packet size. This device allows us to introduce delay, jitter, and a number of different types of error and packet loss on traffic flowing through it. This will cost \$77,020 in capital equipment, and \$11,554 in yearly maintenance for years 2–4. Maintenance can be renewed for a fifth year at the same annual rate. It will interconnect with the FutureGrid core router via 2 10G fibers, to allow 10Gbps in and out of the device.

#### 10.1.4 External Peerings

The FutureGrid network will interconnect with the Internet2 IP network to allow access to developers and users who are not directly connected to FutureGrid. The Indiana GigaPoP will provide this connectivity via its existing Internet2 connections in Atlanta and Chicago. FutureGrid will connect to the Indiana GigaPoP via an existing Indiana GigaPoP 10-Gigabit Ethernet connection between the 710 N Lakeshore Drive facility in Chicago and the Indiana GigaPoP facility on the Indiana University Purdue University Indianapolis campus. There will be no charge for this connectivity from the Indiana GigaPoP or Internet2.

### *10.2 Services Provided by FutureGrid Network*

The FutureGrid network will provide three services.

#### 10.2.1 Isolated Interconnectivity Among Directly Connected FutureGrid Resources

The FutureGrid network's primary service will be to provide interconnection between dedicated FutureGrid resources at the various FutureGrid sites. This will be performed as simply as possible, using simple switching and routing among the sites, and avoiding complex interdomain routing. However, if FutureGrid users require a different configuration, the network may also be



reprovisioned to interconnect sites in other ways, such as using BGP, or at Layer2, making the resources appear to be on the same subnet.

This interconnectivity will be isolated from other networks to allow for more intrusive testing on the network.

#### 10.2.2 Access to Resources Outside of FutureGrid

FutureGrid will also provide, via peering with external networks like Internet2, options for sites outside of FutureGrid to provide resources to FutureGrid, when isolation and dedicated bandwidth are not as important.


#### 10.2.3 Network Impairments

Lastly, the FutureGrid network will allow users to introduce network impairments by selectively routing traffic through a Spirent XGEM network impairment simulator collocated with the FutureGrid core router. This will allow users to introduce jitter, loss, delay, and errors into the network in a fine-grained way using Spirent's built-in TCL interface.

### *10.3 Service Levels*

While FutureGrid is a test-bed environment, it will be crucial that the FutureGrid network perform as expected. The FutureGrid network will be treated as a part of a scientific instrument, providing for availability, repeatability, and transparency. Availability of the FutureGrid network will be vitally important, and any network impairments should be intentional to allow for increased repeatability of tests.

The FutureGrid network will follow standard best practices for maintenance and operations to ensure high availability and predictability for the resource.



## 11. Software Plan

Extensive implementation, adaptation, and, in some cases, extension of open-source software is proposed as part of FutureGrid implementation. Much of that is built on very widely adopted software, such as Torque/MOAB, Xen, MySQL, Pegasus, xCAT, bcfg2, and Inca. In several cases, the FutureGrid team includes original developers of these software packages, and the original developers will make the FutureGrid-required extensions. Core extensions to constituent components will thus be also easily delivered to the larger community outside FutureGrid. We have extensive experience with much of the software that will make up the components of the FutureGrid software environment and have conducted a careful analysis of the software implementation proposed. Further, we have completed preliminary implementations of the Actuating Services software and have developed initial versions of software to instantiate a basic Linux environment, Eucalyptus, and MS HPC Server 2008. The Actuating Services software constitutes one of the most critical elements of the software needed for FutureGrid. The software components of FutureGrid may be subdivided into three basic categories: information services, experiment-management services, and actuating services. These are described in more detail below.

### 11.1 Information Services

We will create a FutureGrid-specific portal using OGCE that will include information and services specific to FutureGrid and FutureGrid use (executed as part of WBS 3.2 by IU and TACC; see Appendix C). These will include information about FutureGrid, details on applying for allocations, access to the Experiment Management and Actuating Services, and access to user data stored within the production components of FutureGrid.

We will integrate into TeraGrid information services and the TeraGrid User Portal (TGUP). We will rely on the TGUP for the following functions:

- *Authentication/Single Sign On (using TeraGrid authentication).* This is a key but potentially confusing point. We will support a variety of authentication mechanisms within FutureGrid, as dictated by user needs. However, the FutureGrid portal, and user data stored within the production elements of FutureGrid (e.g. the Experiment-Management System), will be accessed via the standard TeraGrid authentication system.
- *Help and Information.* We will provide links to general help, information, and documentation about FutureGrid.
- *Allocation Information and allocation management.*

### 11.2 Experiment-Management Services

Experiment management refers to the ability of a test-bed user to define, initiate, and control a repeatable set of events designed to test some particular functionality, either in isolation or in aggregate (executed as part of WBS 3.2 by IU and TACC; see Appendix C). We will deliver the following services via portlets from the OGCE-based FutureGrid portal.

- *Image Browser.* Inspect information about images available for use in FutureGrid.
- *Experiment Browser.* Define experiments as resource, software, and experimental.
- *Software Configuration Browser.* Specify packages and configuration parameters for use in an experiment. Where appropriate, the catalog will indicate any linkages to specific hardware that are dictated by software choices.

- *Monitoring/Instrumentation Browser*. Examine data gathered during experiments.
- *Software Stack Vetting Service*. Submit a software stack for vetting by FutureGrid systems administration staff.
- *Scheduling and Reservation System*. Match researcher requests for test environments against availability.
- *Storage Services*. Store and retrieve all software images and data relevant to a researcher's experiments.


### 11.3 Actuating Services: Test-Bed Management, Experiment Initiation, Collection of Experimental Data, Storage of Data for Later Retrieval and Use

See WBS 3.3 by USC and WBS 3.7 in Appendix C for more information.

- *Authorization Management*. When experiments are scheduled for exclusive use of a test environment in FutureGrid, the first step will be to alter authorization settings so that only the researchers performing that particular experiment may access it.
- *Test Environment Deployment and Configuration*. Install and configure test environments.
- *Experiment Execution*. Experiment execution carries out the steps of an experiment plan.
- *Monitoring and Data Recording*. The monitoring service is responsible for inspecting the state of the test-bed and recording data relevant to the execution of an experiment.

#### 11.3.1 Components of Actuating Services to Implement Dynamically Configurable Environments

The use of Actuating Service software to dynamically configure software environments is one of the most critical elements of FutureGrid. The components of the Actuating Services are well-known, robust, open-source software tools. Their integration is quite novel. The Actuating Services will operate using a suite of software as depicted in Figure 4. See WBS 3.7 in Appendix C for additional information.



**Figure 4: Revised FutureGrid Software Architecture.**

We expect that the open-source core of Eucalyptus will be the primary virtualization environment used for experiments within FutureGrid. We also plan to support experiments with the best-of-breed commercial software including Azure and Microsoft HPC Server.

Inca, OProfile, Vampir, and PAPI will provide system monitoring and instrumentation. As of the writing of this plan, we anticipate using PerfSonar as a tool to collect and record data regarding network conditions.

The scheduling services and system management tools will allow the user environment to be dynamically configurable and support both native OS and virtual OS installations. Eucalyptus and Nimbus will be supported as a layer on top of Xen, which is a core technology required for both software environments. Users will be required to submit contributed environments for vetting prior to installation and availability on dynamic resources. Vetting will be done by Indiana University systems administration staff. This is not likely to be a frequent occurrence, but it is an important capability. Looking backwards, for example, it would have been useful and important to the computational science community to have had this capability to vet and dynamically make available installations of Genesis II.

### 11.3.2 Control of Systems that Are not Dynamically Configurable.

Several systems included within FutureGrid will not be dynamically reconfigurable (see Table 3). Access to these systems will be provided as follows: First, the IU GlobalNOC will enable network paths as appropriate. If configuration changes by hand are needed, these will be put in place by FutureGrid staff. Logins will be enabled for the specific set of user accounts authorized to use one or more of the not-dynamically-configurable systems. Experiments can then proceed.

## **12. Systems Integration and Transition to Operational Status Plan**

In our original proposal, we made the basic assertion that the creation of a highly useful and realistic grid test-bed was a hard problem, and there was no way to create an effective, nationally distributed test-bed without confronting the issues. We thus anticipate spending PY1 in initial deployment of FutureGrid, transitioning at the beginning of PY2 to a fully operational state. This does not mean that development of the facility or software stops at the end of PY1. However, as of the beginning of PY2, we expect FutureGrid to be in a basically operational state.

During all phases of delivery, the systems that comprise FutureGrid will be integrated into current (and future) TeraGrid-wide infrastructure in addition to being integrated into existing infrastructures at IU and partner sites. Corresponding vendors will install systems at all sites and initial hardware diagnostics will be performed.

### *12.1 PY1 – Basic Implementation*

It will be very important to ramp up progress in implementing FutureGrid very quickly, both because of the importance of the project and because of its importance to the evolution of the TeraGrid.

Accelerating work in implementing FutureGrid will involve an acceleration of activities as opposed to a start of activities. The FutureGrid team has had regular teleconferences at least monthly since 3 March 2008. In regular meetings, we have discussed the evolution of plans for the project and worked out several demonstrations in anticipation of a site visit; these are now a basis for demonstrations to be put on at SC09.

Starting in August 2009, each committee will meet on a biweekly basis. As soon as we are in receipt of an award instrument and authorized to begin spending award funds (90 days in advance of award start date, per NSF policy), we will initiate several activities, including

- Weekly meetings of all committees
- Reassignment of existing staff to the FutureGrid project and budget
- Hiring new or reassigning current staff

- Working with the NSF regarding press relations and award announcements
- Finalizing subcontract documents with all subawardees
- Processing purchase orders with vendors, per the purchase contracts

We will hold an all hands face-to-face meeting, likely in Indianapolis, in September. This will be the third such meeting; two were held during the process of developing the proposal. We anticipate a two-day meeting as a way to develop a uniform sense of the PEP; to let various committees meet in parallel and in person; and, most importantly, to energize the team to begin practical work on implementing FutureGrid straight away.

We plan to continue weekly committee calls up to SC09. We anticipate a series of demos and talks in the booths of organizations participating in FutureGrid at SC09. After SC09, we anticipate going to the more measured meeting pace described in Section 4.3.

During this initial implementation year, the systems will be integrated into the local software infrastructure (staff accounts, local customization, file system configuration) at each partner institution. Acceptance testing will begin at this point. Acceptance tests include hardware diagnostics, software functionality testing, performance benchmarking, and stability testing. Acceptance testing criteria is discussed in the vendor contracts. After systems are accepted, they will be integrated into the FutureGrid-wide software infrastructure. This phase will be considered complete when all problems related to hardware and software integration are resolved.

After acceptance and before general operations, systems will be evaluated and hardened for production use. During this transition period, the system will be offered for use to a small group selected from the TeraGrid user community. This transition period will allow for validation of the acceptance test results under a more realistic usage pattern.

## *12.2 PY2 and Beyond – Operational Phase*

Table 3 lists a number of existing systems that will be integrated into FutureGrid, and a number of systems to be acquired during PY1. We will declare FutureGrid to be in an operational mode when three critical elements of the project implementation are all complete: the preexisting systems to be integrated into FutureGrid and systems to be acquired in PY1 are in place and operating as part of FutureGrid; the FutureGrid User Portal is in operation providing basic information and services to FutureGrid users; and users are able to dynamically deploy and use predefined software environments in experiments.

During the operations phase resources will be allocated as follows.

- 10% of resources will be available at the discretion of the PI.
- 90% of resources will be available for allocation to users (who may include members of the FutureGrid team) through a peer-reviewed resource request and allocation process. This process will evolve over time, as detailed below.

As described in section 2.5, the resource request and resource allocation processes will operate in a preliminary learning phase during Program Years 1 and 2, with that process led by co-PI Andrew Grimshaw. We anticipate that effective PY3, it may be both possible and appropriate to transition the process for requesting and awarding resource allocations so that it is somewhat more removed from the FutureGrid team. It seems unlikely that FutureGrid allocations can easily be mixed in with the allocation requests handled by the TRAC, but a process with an external peer review committee may be appropriate. If this seems prudent, then co-PI Grimshaw and

operational staff within TeraGrid will work to create such a panel, assembled using the same principles of peer review as the TRAC and typical NSF peer review processes (merit, no double jeopardy, and guards against real or perceived conflict of interest). One anticipated difference between our vision for a production operation peer review process and the TRAC has to do with format and frequency of meetings. For a resource such as FutureGrid, responsiveness to researchers may require much quicker turnaround than the current TRAC process allows. An operational external peer review process might meet on a monthly basis, and would most likely meet via teleconference rather than in person. It is also the case that mixing and matching resources and requests may be more complicated for FutureGrid than for the TeraGrid. Depending on the particulars, one request may require all or a very large fraction of the FutureGrid resources. Two other requests may require nonoverlapping resources, and so it might be possible to fulfill two different requests simultaneously. As described for PY1 and 2, all projects using time under the PI's 10% discretionary time will be described in a resource request submitted via the same process as any requestor, but submitted as an FYI rather than an action item request.

A critical component of the FutureGrid plan is that we will continue to enhance services over time, particularly the Actuating Services:

- Program Year 2 will add integrated workflows from Pegasus, a storage repository for virtual environments, and scheduling integration.
- Program Year 3 will add instrumentation with Inca and Vampir for the virtual environments.
- Program Year 4 will focus on maintenance of existing technologies and incorporating user needs.

As a result of the pending TeraGrid eXtreme Digital solicitation, there is more uncertainty about future TeraGrid services and processes than usual. We expect FutureGrid to evolve over time, in response to user needs, technological changes, and the plans, processes, and procedures put into place as TeraGrid eXtreme Digital is implemented. We will develop FutureGrid plans and services so as to best meet the needs of the national science and engineering research community in the context of the TeraGrid and the NSF-sponsored national cyberinfrastructure.

## **13. Risk Management Plan**

This risk management plan addresses three primary categories of risks: management risks, operations and facilities risks, and software risks.

### *13.1 Management Risks*

This risk refers to the loss of key members of FutureGrid's management team. The number of institutions involved makes the project more complex but also more resilient; there are more than a half dozen people who could take over overall leadership of the FutureGrid project in the collaboration. The experience of the TeraGrid in general, and IU's early involvement in TeraGrid, revealed a need for considerable depth at the senior management and line management levels. The basic structure of the FutureGrid team includes a PI, four co-PIs, an executive director, a chief architect, and a project manager. This team of eight people takes on a role done in practice at many other TeraGrid participating organizations by two or three people. This team of leaders gives us resilience in the most likely sort of management issue: that some pressing

need related to FutureGrid arises at a time when one person is on vacation or otherwise indisposed.

The Pervasive Technology Institute includes Stewart as executive director and two computer scientists as Center Directors (Geoffrey Fox and Beth Plale). Either Plale or Stewart could fill in for Fox if needed. Similarly, the Research Technologies Division of UIITS, which will manage the system implementation of IU's portion of FutureGrid, now includes four senior managers, any of whom could take over Stewart's current leadership role (Matt Link, D. Scott McCaulay, William Barnett, and Eric Wernert). IU's staffing in research computing and advanced networking has grown significantly since IU was first funded to become part of the TeraGrid, from approximately 100 in 2003 to over 150 today. All in all, the depth and breadth of the FutureGrid team is such that there should be no difficulty in providing proper leadership and management of the project.

### *13.2 Operations and Facilities Risks*

Risks to operational facilities fall into three categories: a) risk that individual hardware will not meet operational requirements; b) risks of physical damage to the housing facilities; and c) risks of scaling the FutureGrid operations across multiple partner sites. Overall, the implementation of hardware should be a very low risk process. The proposed cluster hardware intentionally represents largely standard, oft-used systems so as to make the test-bed as valuable as possible. Given the networking experience of IU and its partners, there is likewise very little technical risk related to networking. Also, the nature of FutureGrid mitigates risks for many users: The use of virtualization will increase portability of applications across partner sites.

#### 13.2.1 Hardware Performance

There is the risk that acceptance testing of hardware will delay operational availability. These delays will be mitigated by the use of a spiral approach to FutureGrid hardware availability and associated capabilities that we offer. This will allow FutureGrid to offer limited capabilities immediately. The initial FutureGrid resource will be the iDataPlex system at IU. This system has already gone through extensive testing using both Eucalyptus and Nimbus and has furthermore been used to prototype our xCAT management system for both Eucalyptus and Nimbus.

Hardware performance is an ongoing risk. First, there is the risk that the initially delivered hardware systems will not meet the vendor-supplied benchmarks. We will mitigate this risk by providing additional cabinets and purchasing additional processors to meet most performance requirements if necessary. There is also the longer-term risk that the processor architecture chosen now will be undesirable in the second half of the project, as many vendors are in a period of transition in their product lines, and GPU-based architectures (for example) could play a larger role than currently forecasted. To mitigate this risk, we have reserved funds for hardware refreshes and upgrades of \$75,000 (PY2), \$250,000 (PY3), and \$75,000 (PY4). We will in general be purchasing systems that are architected to be amendable to future upgrades. Finally, we will maintain vendor support contracts for the duration of the award.

#### 13.2.2 Facilities

The FutureGrid team will prepare disaster recovery plans for all systems and components within FutureGrid. This is standard practice within IU. The Research Technologies Division of UIITS maintains on- and offsite repositories of disaster recovery plans for each service it provides



(currently 169). Purdue has agreed to serve as an alternate site for hosting a cluster should one site become inoperable for an extended period of time.

### 13.2.3 Scaling and Integration of Operations

There is the risk that FutureGrid's distributed operations will be delayed as different members will have different hardware, have different acceptance requirements, and offer different capabilities. Consequently, some of the capabilities of FutureGrid could be delayed. We will mitigate these risks by following a phased approach across system providers, with IU providing FutureGrid's initial systems for users. We will develop a detailed plan for bringing other partners online, but principal details include acceptance testing of hardware, installation of the FutureGrid software stack for managing images, validation of the software, integration of early users onto the site, and finally ramping up to production. Partner sites will be responsible for hardware acceptance. IU will maintain the FutureGrid software stack and will assist partner sites with its installation.

## 13.3 *Software Risks*

The current software plan for FutureGrid can be broken down into two parts: the development of the FutureGrid User Portal and Information Services and the development of the Actuating Services.

Regarding the FutureGrid User Portal and Information services, we believe there is very little technical risk regarding the use of OGCE and other software tools currently used as part of the TeraGrid User Portal (TGUP). This portion of the software development work is simply instantiating a set of standard web components that we have already implemented a number of times. If there is any risk at all, it is mostly risk of available programmer time. This work will be split between IU and TACC, much as is the case for current TGUP development. We believe we have identified the appropriate programmers for this work. If there is a problem in staff availability, we have the option of shifting work between IU and TACC, or from IU and TACC to other FutureGrid participating sites.

As regards Actuating Services, we have two tremendous advantages. First, the actuating services software plan depends upon integration of a number of relatively mature, open-source software products (particularly MOAB and xCAT). Both MOAB and xCAT are mature and well-supported software products that are well suited for our intended purpose. Second, we have already begun pilot implementation of the actuating software services, with good success so far.

We have several risk mitigation options if we run into difficulty with the overall plan or with one or more components of the software that we plan to make up the Actuating Services:

- Our risk-contingency strategy as regards dependence on Pegasus will also be use of scripts. Our risk contingency strategy as regards use of bcfg2 will be use of the open-source image managers Rocks or xCAT.
- Nimbus (a project of FutureGrid partner U. Chicago) can be used to control virtual resource lifecycles similarly to Eucalyptus. However, we anticipate major changes in Nimbus's web service container with the next release of Globus, which is scheduled for release within the first two months of PY1.

- There are additional open-source cloud software projects including the vendor Enomaly (<http://www.enomaly.com/>) and the Open Nebula project (<http://opennebula.org/>) that also provide virtualization software similar to Eucalyptus and Nimbus if necessary.
- We are partners with Grid5000, a similar test-bed project in the EU, and have a mutual commitment for interoperability. This can be pursued more aggressively in our timeline if necessary, and we can adopt a software stack based on what is used within Grid5000. This was considered as our main software plan but rejected because of the different foci of the U.S. and European projects.
- Last but not least, we are already highly experienced and successful at deploying environments through standard Unix scripts. We will develop a set of scripts to instantiate various test environments on all of the systems involved, mitigating risks of our plan to implement automated instantiation of test environments and software stacks.

Software maintenance provides one additional risk: Partners will be responsible for operating software environments on hardware that is part of FutureGrid and owned or operated by other partners. The risk is that the software provider may not have timely access or be available to resolve a problem on another site's hardware. We will mitigate this risk by distributing operational knowledge across multiple sites through training. Our software components, such as Pegasus, are mature and have substantial associated tutorial material. FutureGrid will offer internal tutorials to all partners on the components of our core software stack (described separately).

## 14. Interface Agreements

### 14.1 *FutureGrid Access Mechanisms*

We may divide FutureGrid access mechanisms into two categories: direct access to hardware and access to virtual machines. Direct access to nonvirtualized hardware will use mechanisms familiar from the current TeraGrid:

- GSI-SSH and GSI-SCP: these will use standard TeraGrid accounts and credentials obtained from the TeraGrid MyProxy server. Grid-mapfiles will be updated along with regular UNIX user accounts when a user is granted access to FutureGrid. Grid-mapfile entries will, by default, be the user's NCSA account.
- Regular SSH/SCP: the user has the option of providing a public key to the FutureGrid administrators, who will be responsible for propagating these keys to the user's accounts.
- Community Credentials accessed via MyProxy will also be recognized.
- We will support InCommon login procedures that support account mapping from campus ID to TeraGrid.

FutureGrid's virtualized hardware will be accessed via SSH keys generated by the appropriate cloud technology (Eucalyptus, by default). Eucalyptus site administrators will issue keys. Nimbus-based access to virtual machines will use TeraGrid MyProxy keys.

In addition to these basic, bootstrap mechanisms, FutureGrid by its nature may need to support additional authentication mechanisms. These include both command-line tools and web-based methods. Command-line tools include GSSAPI methods for single sign on. Web methods include CAS, Shibboleth, SAML, OpenID, and OAuth. The later web methods are appropriate for providing access to web sites and web services that may be offered by a testing application.

These are also often associated with social networking applications, and we will leverage the expertise and ViNe technology of our UF partners. Any additional methods will be evaluated for level of maturity as well as benefits and pitfalls before inclusion in the suite of possible access mechanisms.

We will develop an extensive set of use cases based on our first year's outreach activities.

### *14.2 Networking*

All FutureGrid sites with dedicated network connections will be responsible for providing networking equipment at their institution to connect the resource(s) at their site to the FutureGrid backbone network.

FutureGrid bandwidth will be available for reservation. Requests for bandwidth should be provided at least a week prior to when the reservation is requested. Requests will routinely be made via email to the normal FutureGrid help address, from whence they will be forwarded to staff of the IU GlobalNOC. Staff of the GlobalNOC will be responsible for fulfilling these requests. When specific bandwidth or network impairment is requested between one or more FutureGrid systems, those systems will be restricted to only the requested application using access control lists and firewalls.

### *14.3 Security*

Protocols for authentication and authorization will follow current TeraGrid standards (SSH Keys/X.509).

In the event of critical security issues, the TeraGrid incident response team will be alerted as described in the TeraGrid Security Working Group Security Playbook. Sites will report incidents using the TeraGrid Security Incident Response Form. Each FutureGrid site will have a designated point of contact for security who will coordinate communications between the FutureGrid sites, TeraGrid incident response, and secondary responders. All communications between responders must be encrypted via PGP/GPG. As soon as possible, and within 24 hours of incident discovery, incidents will be reported to the parties listed in section 15.3.

### *14.4 Accounts*

Responsibility for user accounts will rest on the system administrator already responsible for TeraGrid accounts at each site that is already a TeraGrid resource provider. For FutureGrid sites that are not current TeraGrid RPs, one system administrator and one backup person will be designated as having responsibility for user account creation and maintenance.

Account information will be propagated between sites using the standard system used by the TeraGrid, or its successor software. This is currently AMIE, a system created by the TeraGrid.

In the event account suspension is required, suspension of accounts will be handled via the same channels the TeraGrid uses (AMIE, coordination with local system administrators). Account suspension due to a security incident will occur within 24 hours of notification. Suspension or general account requests not security related will be handled next business day.

FutureGrid partners will be responsible for implementing the policies described by the TeraGrid for handling accounts at individual sites. Enforcement of the policies will be the responsibility of

Dave Hancock at IU, backed by a hardware committee made up of representatives from each FutureGrid site.

### *14.5 Services*

Information about the services available at each FutureGrid site will be published on a FutureGrid central web site, as well as on the TeraGrid User Portal and via announcements to the TeraGrid community.

Each FutureGrid site is expected to provide the following services on their resource:

- GSISSH
- GridFTP
- Modules as specified by the TeraGrid (the transition to Modules is planned for Spring 2010).
- Interoperability with xCAT and bcfg2
- Lustre-compatible kernel for client systems
- Xen-compatible kernel for compute nodes for use by Eucalyptus or Nimbus
- Inca
- PAPI or similar hardware performance environment
- Vampir
- Resource management and scheduling integration

#### 14.5.1 Software Support and Deployment

Responsibility for support and deployment of software services provided by each site is as follows.

IU: MOAB, Torque, Bcfg2, xCAT, PerfSonar, Eucalyptus

USC: Pegasus

TACC: Test harness

UC/ANL: Nimbus, CTSS

UCSD: Inca

UF: ViNe, network appliance, Social VPN

UTK: PAPI

UV: Genesis II, Unicore, g-lite

Technische Universitaet Dresden: Vampir development, easy access to prerelease versions of Vampir and VampirTrace

GWT-TUD GmbH: Vampir support

Software packages and images will be deployed using IU's Lustre WAN file system as a central repository with SCP as a backup mechanism to propagate central software to all sites. Sites will be notified of a new software release and installation will follow FutureGrid change management procedures overseen by the FutureGrid Project Manager.

User software requests will be accepted and responses are expected within 5 working days.

We will group test environments into three categories in terms of levels of support: those for

which the FutureGrid team offers extensive support and debugging of applications (TeraGrid CTSS, Eucalyptus, Nimbus, Genesis II); those for which the FutureGrid team offers some consulting support, but will also depend upon established communities of users or corporate providers for support (e.g., Windows HPC Server, Xen, VMware, EGEE/g-lite; Unicore, Condor, BOINC); and user-provided test environments. Those who provide their own test environments will be expected to be self-supporting. Depending upon patterns of usage over time, levels of support will be adjusted to match FutureGrid researcher needs.

Test environments will be instantiated in accordance with researcher needs and allocations of time on FutureGrid resources. We expect the ambient (default) state of FutureGrid systems to be as follows: The high throughput cluster at Purdue will run Condor and the CTSS. The Cray XT5m and a shared memory system to be identified will run the vendor-recommended Linux OS and the current version of CTSS. The IBM iDataPlex systems at UF and UC/ANL will run Nimbus. The IBM iDataPlex at IU will have a small number of nodes dedicated to Microsoft Windows HPC Server. The remainder of the iDataPlex at IU, the iDataPlex at UCSD/SDSC, and the Dell PowerEdge at TACC will be dynamically reconfigurable.

#### 14.5.2 Data Flow and Storage

##### *14.5.2.1 Data Storage Systems*

###### IU

Statically configured 339TB Lustre WAN file system. The Lustre-WAN file system is theoretically capable of sustained I/O of 3.2 GBps to locally connected FutureGrid systems. Connections from the remainder of the FutureGrid systems are limited by the bandwidth of the connection to IU of 10 Gbps.

Statically configured 6TB HPSS test instance, 2.8PB HPSS production instance for experiment data

###### TACC

Statically configured 15TB NFS file system

###### UC/ANL

Statically configured 120TB GPFS file system

###### UCSD/SDSC

Statically configured 72TB Lustre or PVFS file system

###### IU, TACC, UC, UCSD, UF

Each of the dynamically configurable resources will have at least eight configurable storage nodes with at least 16TB of aggregate attached storage.

##### *14.5.2.2 Data Flow*

These storage systems will form a hierarchy for the dynamically configurable computational resources. Data flow will revolve around the 335TB Lustre WAN file system at IU. Sites will use their statically configured storage as a local cache of images that are transferred to and from the Lustre WAN file system. Locally at each site, images will be passed from the local cache to the

nodes within a cluster for instantiation. After modification of images or experiment data, they will be transferred back to the Lustre WAN file system and archived within HPSS if desired.

Data transfers will initially rely on a combination of Lustre WAN file system mounts from IU at each site and GridFTP transfers between resources. This model will support existing data transfer tools integrated within the TeraGrid User Portal. Data storage to and from dynamic resources will be handled by Pegasus, including archiving to HPSS, after Pegasus is fully integrated into FutureGrid.

### *14.6 Operations*

All FutureGrid sites are coordinating their activities across sites, which are expected to supply 24x7 availability with set preventative maintenance windows.

In case of severe security vulnerabilities or system issues that impact the availability of the system, emergency maintenance will be undertaken in order to correct the issue.

Both preventative maintenance and outages will be communicated via the [news.teragrid.org](http://news.teragrid.org) information system.

FutureGrid systems will be monitored at the IU Global Research Network Operations Center.

### *14.7 Support*

We will adopt the tiered support model, championed for many years. This model sets support into the following tiers, with support generally moving from first to last.

- **Online help accessed 24x7 by users, via the IU/TeraGrid Knowledge Base.** A very large fraction of user questions are successfully resolved by providing excellent documentation in an easily searchable format. Problems that are not resolvable through accessing the Knowledge Base may be escalated immediately to telephone support or, on a slightly longer timeline, by filling out a web form or sending email.
- **24x7 phone support.** IU will provide 24x7 phone support delivered from the Global Research Network Operations Center (GRNOC). The GRNOC will provide 24x7 system status information, immediate handling of security concerns and incidents, and limited technical support, and will either forward phone calls to second-level technical experts (between 8 am and 8 pm Eastern Time) or initiate a trouble ticket (between 8 pm and 8 am Eastern Time).
- **Second-tier telephone and email support.** Technical experts at IU will provide second-tier support to users via telephone or email (in response to email or web form queries). For some systems problems, it may regularly be the case that second-tier problems are referred to systems management personnel at sites hosting FutureGrid hardware when a problem appears to be specific to a particular machine.
- **Third-tier support.** Top technical experts anywhere within the FutureGrid team will provide third-tier support. Such experts may also be involved in advanced user support provided via the TeraGrid or TeraGrid XD.

Throughout the tiered user support/problem resolution process, we will use a trouble tracking system to ensure that user issues are promptly addressed. We will also ensure that at all times users have a single point of contact within FutureGrid and know who that point of contact is.

We will use Web 2.0 services to allow users to share experiences, and to enable one-to-many and many-to-many discussions in resolving problems and enabling new capabilities. For a resource that serves a community of leading grid experts, enabling users to share expertise should be particularly beneficial.

Personnel supporting software and applications that execute on another site's hardware will be provided privileged access in accordance with best practices for each operating system using the principle of least privilege (<http://hissa.ncsl.nist.gov/rbac/paper/node5.html>). This will decrease time to problem resolution and allow the FutureGrid software environment to be supported by the original developers of the software in most cases. Second- and third-tier support personnel will be provided training on common FutureGrid software environments to develop local experts at each organization.

### *14.8 Management*

The FutureGrid management structure will be as described in section 4. Each organization participating in FutureGrid will participate in committees relevant to that organization's involvement. (For example, it is not expected that USC, which is participating strictly via involvement in FutureGrid software, will participate in the Systems Administration and Network Management Committee.) Every participating organization will participate in the Operations and Change Management Committee. Each site will have a primary and secondary representative for project management and reporting.

Most committees will meet on a monthly basis, as described in section 4.3. The Operations and Change Management Committee will have one major meeting monthly, and weekly status check/change management meetings.

Decisions of FutureGrid committees and leadership management will be transmitted via email. Email will be sent to each participating organization's members of the Operations and Change Management Committee, each site's primary and secondary representative for project management and reporting, each site's lead investigator, and all members of the committee making a decision. There is a conflict escalation process described in the Management section 4.4. Abiding by and enacting decisions made by the relevant committees, or settled through the conflict escalation process, will be a contractual term of participating institution subcontracts.

For each award and subaward site, the sponsored research office for that site will carry out the necessary agreements for the award to be accepted. Turnaround time for each sponsored research office will vary, but it is estimated that the first subaward will be complete within 4 to 6 weeks. Subcontractors will all comply with the terms specified by Indiana University and detailed in Appendix F.

## **15. Cybersecurity Plan**

Cybersecurity will be integral to FutureGrid and affects three main areas: computer systems, data, and software. FutureGrid will follow standard best practices to ensure that its systems are not vulnerable to cyber attacks.

## *15.1 System and Infrastructure Security*

### *15.1.1 Infrastructure Security*

#### *15.1.1.1 Indiana University*

Access to IU facilities (Innovation Center, Wrubel Computing Center, Cyberinfrastructure Building) will be controlled by key-card access to offices on a granular basis determined by the requirements of staff roles. Access to Data Center facilities will be restricted to system administrators who will be performing physical maintenance on machines, for electrical or network maintenance. Fire suppression will be provided by a double interlock system and accompanied with a Very Early Smoke Detection Apparatus (VESDA).

The IU Data Center facility is monitored 24 hours a day, 7 days a week. Operations staff monitor the facility by CCTV, and the card key system records accesses to rooms by person. During evenings and weekends, IU Police Department officers are present at the facility.

#### *15.1.1.2 Purdue University*

Centralized computer facilities that house core data will be protected in a physically secure location with controlled access. Computer facilities that process departmental data may require physical security depending on the value and sensitivity of the data they process, the resources they access, and their cost.

Fire suppression is provided accordance with Purdue University standards and FM Global requirements. The center itself is protected by a dry-pipe, double-interlocked preaction sprinkler system following university risk-management policy. This system is tested following Purdue University and state of Indiana standards.

#### *15.1.1.3 TACC*

Physical security of the TACC facility is ensured through several measures. The machines are secured via a card-key access system limited to TACC staff only and monitored by a 24x7 operations staff. TACC User Services Staff are able to view the system through large glass windows installed on two sides of the machine room to guard against unauthorized access.

TACC has a shutdown procedure plan for both emergency and nonemergency situations. There are system control alarms for systems when either temperature at various points in the room exceeds a certain threshold or if the flow of chilled water should be interrupted. In addition, during off hours, Operations staff walk the machine room floor every 4 hours to detect environmental issues.

The fire detection system for the TACC Commons computer room is separate from that of the main building. The detection system is configured in two zones and is tied to a Halon 1301 fire suppression system. Fire conditions in either zone will initiate alarms, but fire conditions in both zones must exist to initiate a Halon dump. There are manual pull stations at multiple locations that will place the system into alarm and initiate a dump.

Finally, the TACC office building where the machine room is located is contained on a fenced and gated secure facility at the J.J. Pickle Research Campus. UT security guards monitor the gates and check identification outside of normal business hours to confirm that individuals have previous authorization to enter campus grounds after hours.



#### *15.1.1.4 UC/ANL*

Physical access to UC resources such as servers, storage, switches, racks, and firewalls is restricted to staff and certain others who have a need for such access and have been granted prior authorization by designated Networking Services and Information Technologies (NSIT) management. NSIT is able to audit physical access. Non-NSIT contractors, e.g. vendor service personnel, who require physical access will either be given a time bounded access code or token which will grant them access or they will be granted one-time access by authorized Data Center Operations staff. Fire/smoke detectors and preaction water sprinklers provide fire suppression with an automatic emergency electrical shutoff.

#### *15.1.1.5 UCSD*

The SDSC Datacenter is secured with biometric access controls, 24x7 operations staff, and video surveillance. Only staff directly involved in administration of machines in the datacenter, limited management personnel, and guests accompanied by authorized personnel have access to the machine room. Fire suppression is provided by Halon with a two-stage (dry-pipe) water backup system. Two independent electrical circuits, either of which can fully power the building, power the datacenter. UCSD also operates its own cogeneration facility capable of supplying campus loads even if regional grid power should fail. Critical systems equipment is protected from power events via Uninterruptible Power Supplies (UPS). The average number of scheduled outages per year for electrical and cooling maintenance is less than 1, with an average annual impact of 6 hours of planned outages related to facilities management.

#### *15.1.1.6 UF*

The ACIS lab machine room has secured code-key access and daytime video surveillance/recording. Only staff directly involved in systems or network administration have access to the machine room. Fire and smoke detectors and water sprinklers provide fire suppression. FutureGrid equipment will be protected from power failures via short-time UPS backup. The total annual number of scheduled outages per year for electrical and cooling maintenance is planned to be two, for a total of 48 hours of planned outages related to facilities management per year.

### 15.1.2 System Security

All systems are to be maintained with a standard maintenance schedule (first Tuesday of every month).

Critical (remote-root exploit) vulnerabilities may be addressed by emergency maintenance, with notification through proper channels both locally on site as well as to the other grant sites via mailing list or RSS feed.

Noncritical vulnerabilities are to be addressed during standard maintenance.

All maintenance windows are to be announced through proper channels both locally on site as well as to the other grant sites via mailing list or RSS feed.

In case of an availability-based attack, administrators will work with on-site administrators of other services (the Global Research NOC, REN-ISAC, University Information Security Office) to filter or block traffic from attacking sites.

Where applicable, compute systems will be restricted to a private network; only head/submit nodes will be available through the public network.

Any nodes available via the public network will have host-based firewalls and other access-control methods installed, including one-time passwords for administrative users.

Access to systems will be via encrypted channels (e.g. ssh for maintenance of the system, https for web services).

Regular backups will be made of critical machines in case of hardware failure. We plan for weekly full backups and daily incremental backups.

## *15.2 Software Security*

Software developed for testing on FutureGrid hardware will follow standard best practices for user authentication and authorization

Code developed for use on FutureGrid hardware will be version controlled and make use of automated testing.

Software from other sources installed on FutureGrid hardware will be examined for security vulnerabilities. Administrators will follow software project news in order to be aware of security vulnerabilities as they are discovered. In case of security vulnerabilities in software, maintenance to address the issue will be taken based on the severity of the issue as detailed above.

Any service running on FutureGrid hardware will be tested for its ability to handle malformed or incorrect inputs (black box testing) as well as its ability to handle malicious attacks such as SQL injection or buffer overflows (white box testing).

Data produced and used for computation on the FutureGrid hardware will be likely be test data, but all necessary steps will be taken to securely destroy data. Generated data is destroyed after usage automatically to guarantee confidentiality and trust in the system. Configuration, authentication, or other sensitive information will be transmitted via encrypted channels. Information on FutureGrid hardware that is discarded at the end of its lifecycle will be removed via standard procedures for data destruction.

## *15.3 Incident Response*

Each FutureGrid site will have a primary and secondary designate for handling security incidents. At FutureGrid participant sites that are currently part of the TeraGrid, these designates will be the existing designates for reporting TeraGrid security incidents. At FutureGrid sites that are not currently TeraGrid sites, primary and secondary security incident handlers will be designated following guidelines established by the TeraGrid for incident response.

Severe security incidents will be immediately reported to the FutureGrid PI and co-PIs, IU Director of Research Technologies Systems Matt Link and the Research Technologies Systems management team, and the IU University Information Security Office, as well as administration at other FutureGrid sites and the NSF Program Office. This plan follows the model of escalation for TeraGrid security incidents outlined by TeraGrid's security working group. Severe incidents include unauthorized root/administrator access to FutureGrid machines, resource-based attacks (denial of service, widespread virus/malware, botnet attack, etc.), and attacks on mission-critical applications or servers.

Nonurgent security incidents (e.g. unsuccessful attempts at severe attacks, degradation of service attacks) will be reported within 24 hours to the same people listed above via encrypted email and a centralized incident response tracking system. Electronic incident response will use the same policies as incident response for the TeraGrid, but the FutureGrid mailing list and electronic incident response will be separate systems from the TeraGrid's existing infrastructure.

Incident response will be handled via an issue-tracking application, to which all of the above individuals will have access and which will record the reporting information detailed below.

Administrators will provide a report with as much detail as possible on the security incident, including

Date and time incident was detected

Date and time incident actually occurred (if different from above)

Type of incident (e.g., web defacement, virus/worm, etc.)

Method of intrusion (e.g., vulnerability exploited), if known

Level of unauthorized access attained (e.g., root, administrator, user, etc.), if known

Log extracts (if appropriate and available)

## Appendix A: FutureGrid Project Plan Milestone Schedule

| WBS     | Milestone   | YR | Finish | Metric  |
|---------|---|----|--------|---|
| 1.0     | <b>Hardware</b>   |    |        |   |
| 1.1.4.4 | Dell 1152 core hardware installation completed                          | 1  | Jan-10 | TACC cluster ready for acceptance testing                           |
| 1.1.5.1 | Dell 1152 core hardware acceptance test completed                       | 1  | Mar-10 | Configurations meet or exceed those proposed by vendors in contract |
| 1.1.5.3 | Initial Dell 1152 core hardware cluster completed                       | 1  | Apr-10 | TACC cluster available for use                                      |
|         |   |    |        |   |
| 1.2.4.4 | IBM iDataPlex 1024 core hardware installation completed                 | 1  | Dec-09 | IU IBM cluster ready for acceptance testing                         |
| 1.2.5.1 | IBM iDataPlex 1024 core acceptance test completed                       | 1  | Mar-10 | Configurations meet or exceed those proposed by vendors in contract |
| 1.2.5.3 | Initial IBM iDataPlex 1024 core completed                               | 1  | Apr-10 | IU IBM cluster available for use                                    |
|         |   |    |        |   |
| 1.3.4.4 | IBM iDataPlex 672 core hardware installation completed                  | 1  | Dec-09 | UC IBM cluster ready for acceptance testing                         |
| 1.3.5.1 | IBM iDataPlex 672 core acceptance test completed                        | 1  | Mar-10 | Configurations meet or exceed those proposed by vendors in contract |
| 1.3.5.3 | Initial IBM iDataPlex 672 core completed                                | 1  | Apr-10 | UC IBM cluster available for use                                    |
|         |   |    |        |   |
| 1.4.4.4 | IBM iDataPlex 256 core hardware installation completed                  | 1  | Dec-09 | UF IBM cluster ready for acceptance testing                         |
| 1.4.5.1 | IBM iDataPlex 256 core acceptance test completed                        | 1  | Mar-10 | Configurations meet or exceed those proposed by vendors in contract |
| 1.4.5.3 | Initial IBM iDataPlex 256 core completed                                | 1  | Apr-10 | UF IBM cluster available for use                                    |
|         |   |    |        |   |
| 1.5.4.4 | Cray XT5M 672 core hardware installation completed                      | 1  | Dec-09 | IU Cray cluster ready for acceptance testing                        |
| 1.5.5.1 | Cray XT5M 672 core acceptance test completed                            | 1  | Mar-10 | Configurations meet or exceed those proposed by vendors in contract |
| 1.5.5.3 | Initial Cray XT5M 672 core completed                                    | 1  | Apr-10 | IU Cray cluster available for use                                   |
|         |   |    |        |   |
| 1.6.4   | Shared memory cluster acquisition completed                             | 2  | Nov-10 | Cluster ready for acceptance testing                                |
| 1.6.5.1 | Shared memory cluster acceptance test completed                         | 2  | Dec-10 | Configurations meet or exceed those proposed by vendors in contract |
| 1.6.5.4 | Shared Memory cluster completed   | 2  | Apr-11 | Shared memory cluster ready production use                          |
|         |   |    |        |   |
| 1.7.3.3 | IBM iDataPlex 256 core hardware installation completed                  | 1  | Jan-10 | UCSD IBM cluster ready for acceptance testing                       |
| 1.7.4.1 | IBM iDataPlex 256 core acceptance test completed                        | 1  | Feb-10 | Configurations meet or exceed those provided by IU                  |
| 1.7.4.3 | Initial UCSD IBM iDataPlex 672 core completed                           | 1  | Mar-10 | UCSD IBM cluster available for use                                  |
|         |   |    |        |   |
| 1.9.3   | DataDirect Networks S2A6620 Storage Appliance acquisition completed     |    | Oct-09 | Storage Appliance ready for acceptance testing                      |
| 1.9.4.1 | DataDirect Networks S2A6620 Storage Appliance acceptance test completed |    | Nov-09 | Configurations meet or exceed those proposed by vendors in contract |
| 1.9.4.2 | DataDirect Networks S2A6620 Storage Appliance completed                 |    | Dec-09 | Storage Appliance ready for production use                          |
|         |   |    |        |   |
| 1.10.3  | Sun X4540 Storage Server acquisition completed                          |    | Oct-09 | Storage servers ready for acceptance testing                        |

| WBS       | Milestone  | YR | Finish | Metric  |
|-----------|--|----|--------|---|
| 1.10.4.1  | Sun X4540 Storage Server acceptance test completed |    | Nov-09 | Configurations meet or exceed those proposed by vendors in contract   |
| 1.10.4.2  | Sun X4540 Storage Server completed                 |    | Dec-09 | Storage servers ready for production use  |
| 2.0       | <b>Networks</b>                                    |    |        |   |
| 2.1       | Network contracts completed                        | 1  | Oct-09 | Signed contracts in Purchasing  |
|           | CENIC  |    |        |   |
|           | Starlight  |    |        |   |
|           | FLR  |    |        |   |
|           | NLR  |    |        |   |
|           | AT&T   |    |        |   |
|           | Spirent  |    |        |   |
|           | Matrix Integration                                 |    |        |   |
| 2.2       | Core router installation completed                 | 1  | Dec-09 | Connectivity to sites measured by bandwidth between sites   |
| 2.3       | Network impairments simulator installed            | 1  | Dec-09 | Programmatic introduction of network latency, jitter, loss, and errors available                              |
| 2.4       | IU, UC, UF, UCSD, and TACC connectivity completed  | 1  | Feb-10 | Network to Chicago working  |
| 2.6       | TeraGrid and Internet2 connectivity completed      | 1  | Mar-10 | Network to TeraGrid and Internet2 working   |
| 3.0       | <b>Software</b>                                    |    |        |   |
| 3.1       | AMIE   | 1  | Sep-11 | Reporting data to TeraGrid  |
| 3.2       | User Portal  |    |        |   |
| 3.2.1.1   | Portal design completed                            | 1  | Dec-09 | Portal design document available for review   |
| 3.2.1.2   | Authentication/single sign                         | 1  | Mar-10 | Beta version of user portal available for use   |
| 3.2.1.3   | Portal resource availability tracking completed    | 1  | Jul-10 | Resource data available in portal   |
| 3.2.1.4   | Links to general help information completed        |    | Jul-10 | Links to general help, information, and documentation about FutureGrid successfully tested                    |
| 3.2.1.5   | Initial version of User Portal completed           | 2  | Sep-10 | User Portal ready for production use  |
| 3.2.2.1.3 | Image Browser deployed                             | 2  | Jan-11 |   |
| 3.2.2.2.3 | Experiment Browser deployed                        | 2  | Feb-11 |   |
| 3.2.2.3.3 | Software Configuration Browser deployed            | 2  | Mar-11 |   |
| 3.2.2.4.3 | Monitoring/Instrumentation Browser deployed        | 2  | Mar-11 |   |
| 3.2.2.5.3 | Scheduling, reservations deployed                  | 3  | Dec-11 | Provide capability of matching researcher requests for test environments against availability                 |
| 3.2.2.6.3 | Storage services deployed                          | 2  | Sep-11 | Provide capability to store and retrieve all software images and data relevant to a researcher's experiments. |
| 3.2.2     | Experiment   | 2  | Jan-11 | Portal interface to view/manage user/group information available for use                                      |
| 3.2.3     | Portal user information management completed       | 2  | Jan-11 | Portal interface to view/manage user/group information available for use                                      |
| 3.2.4     | Test harness access completed                      | 2  | May-11 | Test harness accessible via portal  |
| 3.2.5     | Portal maintenance – PY2 H1                        | 2  | Mar-11 | Portal updated  |
| 3.2.6     | Portal maintenance – PY2 H2                        | 2  | Sep-11 | Portal updated  |
| 3.2.7     | Portal maintenance – PY3 H1                        | 3  | Apr-12 | Portal updated  |
| 3.2.8     | Portal maintenance – PY3 H2                        | 3  | Sep-12 | Portal updated  |
| 3.2.9     | Portal maintenance – PY4 H1                        | 4  | Mar-13 | Portal updated  |

# FutureGrid: An Experimental, High-Performance Grid Test-Bed

| WBS     | Milestone   | YR | Finish | Metric   |
|---------|---|----|--------|--|
| 3.2.10  | Portal maintenance – PY4 H2   | 4  | Sep-13 | Portal updated   |
| 3.3     | Pegasus   |    |        |  |
| 3.3.1   | Pegasus available on test-bed   | 1  | Mar-10 | Pegasus available for use  |
| 3.3.2   | Pegasus documentation completed   | 1  | Mar-10 | Available on project web site  |
| 3.3.3   | Pegasus immediate resource provisioning workflow completed              | 1  | Jun-10 | Resource provisioning available for use  |
| 3.3.4   | Pegasus time-sensitive resource provisioning workflow completed         | 2  | Jun-11 | New time-driven tasks available in Pegasus workflows   |
| 3.3.5   | Workflow repository requirements completed                              | 2  | Sep-11 | Requirements documented for development  |
| 3.3.6   | Pegasus tutorial completed  | 2  | Dec-10 | Available on project web site  |
| 3.3.7   | New end-to-end workflows added to Pegasus                               | 3  | Sep-12 | End-to-end workflows from resource provisioning to injection of events available                         |
| 3.3.8   | Pegasus workflow repository completed                                   | 4  | Mar-13 | Web access to workflow repository available  |
| 3.4     | Grid Benchmark Challenge  |    |        |  |
| 3.4.1   | PAPI supported at all FutureGrid sites                                  | 2  | Dec-10 | Ability to measure low-level performance on all FutureGrid computers                                     |
| 3.4.2   | HPCC benchmark with Globus/MPICH-G                                      | 2  | Mar-11 | HPCC benchmark works with Globus/MPICH-G   |
| 3.4.3   | Modifications of HPCC network tests for cross-site execution completed  | 2  | Sep-11 | GBC has cross-site network component   |
| 3.4.4   | Modifications of local computational tests of HPCC benchmark completed  | 3  | Dec-11 | GBC runs local tests across sites to track variability in hardware speeds                                |
| 3.4.5   | Modifications of global computational tests of HPCC benchmark completed | 3  | Jun-12 | GBC runs gobal tests across sites  |
| 3.4.6   | Virtualization of HPCC benchmark completed                              | 4  | Mar-13 | GBC works on virtualizaed hardware   |
| 3.4.7   | Heterogeneous virtualization of HPCC benchmark completed                | 4  | Jun-13 | GBC works in a mixed virtual environment   |
| 3.5     | Inca  |    |        |  |
| 3.5.1.1 | Testing plan for monitoring FutureGrid functionality completed          | 1  | Dec-09 | Document of planned and needed tests available on project web site                                       |
| 3.5.1.2 | Initial Inca functionality deployment complete                          | 1  | Dec-09 | Grid monitoring in use   |
| 3.5.1.3 | User documentation complete   | 1  | Dec-09 | Available on project web site  |
| 3.5.1.4 | Add additional tests as new software is added or updated                | 1  | Sep-10 | Inca upgraded  |
| 3.5.1.5 | Support NSF required and optional benchmarks                            | 1  | Sep-10 | Inca upgraded  |
| 3.5.2   | Add additional tests/benchmarks YR2                                     | 2  | Sep-11 | Inca upgraded  |
| 3.5.3   | Add additional tests/benchmarks YR3                                     | 3  | Sep-12 | Inca upgraded  |
| 3.5.4   | Add additional tests/benchmarks YR4                                     | 4  | Sep-13 | Inca upgraded  |
| 3.6     | Nimbus  |    |        |  |
| 3.6.1   | Nimbus deployments completed  | 1  | Sep-10 | Nimbus available for uses on all FutureGrid clusters   |
| 3.6.2   | Nimbus maintenance – PY2  | 2  | Sep-11 | Nimbus upgraded  |
| 3.6.3   | Nimbus maintenance – PY3  | 3  | Sep-12 | Nimbus upgraded  |
| 3.6.4   | Nimbus maintenance – PY4  | 4  | Sep-13 | Nimbus upgraded  |
| 3.7     | Actuating Services  |    |        |  |
| 3.7.1.4 | FutureGrid components completed   | 1  | Apr-10 | All components necessary to support instantiation of virtual and real machines successfully implemented, |
| 3.7.2.1 | Xen instantiation completed   | 2  | Apr-10 | Xen virtual monitor in production  |

| WBS     | Milestone   | YR | Finish | Metric  |
|---------|---|----|--------|---|
| 3.7.2.2 | Eucalyptus instantiation completed                                  | 2  | Oct-11 | Ability to instantiate virtual machines via Eucalyptus supported      |
| 3.7.2.3 | Nimbus instantiation completed                                      | 2  | Oct-11 | Ability to instantiate virtual machines via Nimbus supported          |
| 3.7.2.4 | VMWare instantiation completed                                      | 2  | Oct-11 | Ability to instantiate virtual machines via VMWare supported          |
| 3.7.2.5 | RPM-based Linux instantiation completed                             | 2  | Oct-11 | Ability to instantiate RPM-based Linux machine supported              |
| 3.7.2.6 | Windows 2008 instantiation completed                                | 2  | Oct-11 | Ability to instantiate machine running Windows 2008 supported         |
| 3.7.2.7 | Microsoft HPC Server instantiation completed                        | 2  | Oct-11 | Ability to instantiate machine running Microsoft HPC Server supported |
| 3.8     | ViNe  |    |        |   |
| 3.8.1   | ViNe routing API and middleware completed                           | 1  | Sep-10 | ViNe integrated with FutureGrid                                       |
| 3.8.2   | ViNe management interfaces completed                                | 2  | Sep-11 | ViNe upgraded   |
| 3.8.3   | ViNe management services completed                                  | 3  | Sep-12 | Programmatic ViNe management APIs available                           |
| 3.8.4   | ViNe refactoring and improvements completed                         | 4  | Sep-13 | ViNe upgraded   |
| 3.9     | SocialVPN   |    |        |   |
| 3.9.1.4 | VM and Facebook versions of SocialVPN virtual appliance completed   | 1  | Mar-10 | Number of virtual appliance downloads; number of deployed appliances  |
| 3.9.1.5 | On-line tutorial and video completed                                | 1  | Mar-10 | Number of social network users; web page and video hits               |
| 3.9.1.6 | Social network bindings completed                                   | 1  | Sep-10 | OpenSocial, Skype, and Gmail chat                                     |
| 3.9.2   | Education modules and updated tutorial/video completed              | 2  | Sep-11 | Number of virtual appliance downloads; number of deployed appliances  |
| 3.9.3   | Virtual appliance enhancements and updated tutorial/video completed | 3  | Sep-12 | Number of virtual appliance downloads; number of deployed appliances  |
| 3.9.4   | Virtual appliance enhancements and updated tutorial/video completed | 4  | Sep-13 | Number of virtual appliance downloads; number of deployed appliances  |
| 3.10    | Test Harness  |    |        |   |
| 3.10.1  | Initial test harness with limited functionality completed           | 2  | Sep-10 | File transfers; start/stop agents; command-line interface             |
| 3.10.2  | Test harness logging completed                                      | 2  | Mar-11 | Merging distributed logs into a unified experiment log                |
| 3.10.3  | Web interface completed   | 2  | Sep-11 | Web interface for managing experiments available                      |
| 3.10.4  | Test harness maintenance – PY2 H1                                   | 2  | Mar-11 | Test harness upgraded   |
| 3.10.5  | Test harness maintenance – PY2 H2                                   | 2  | Sep-11 | Test harness upgraded   |
| 3.10.6  | Test harness maintenance - PY3 H1                                   | 3  | Mar-12 | Test harness upgraded   |
| 3.10.7  | Test harness maintenance – PY3 H2                                   | 3  | Sep-12 | Test harness upgraded   |
| 3.10.8  | Test harness maintenance – PY4 H1                                   | 4  | Mar-13 | Test harness upgraded   |
| 3.10.9  | Test harness maintenance – PY4 H2                                   | 4  | Sep-13 | Test harness upgraded   |
| 3.11    | Genesis II, Unicore, and EGEE                                       |    |        |   |
| 3.11.1  | Genesis II, Unicore, and EGEE deployments completed                 | 1  | Sep-10 | Software running on all FutureGrid nodes                              |
| 3.11.2  | Genesis II, Unicore, and EGEE maintenance YR2                       | 2  | Sep-11 | Genesis II, Unicore, and EGEE upgraded                                |
| 3.11.3  | Genesis II, Unicore, and EGEE maintenance YR3                       | 3  | Sep-12 | Genesis II, Unicore, and EGEE upgraded                                |

# FutureGrid: An Experimental, High-Performance Grid Test-Bed

| WBS     | Milestone   | YR | Finish | Metric  |
|---------|---|----|--------|---|
| 3.11.4  | Genesis II, Unicore, and EGEE maintenance YR4               | 4  | Sep-13 | Genesis II, Unicore, and EGEE upgraded                                      |
| 3.12    | Vampir  |    |        |   |
| 3.12.1  | VampirServer deployment completed                           | 1  | Sep-10 | Software running on central server at IU                                    |
| 3.12.2  | VampirTrace deployments completed                           | 1  | Sep-10 | Software running on all FutureGrid nodes                                    |
| 3.12.3  | Vampir maintenance YR2                                      | 2  | Sep-11 | Vampir upgraded   |
| 3.12.4  | Vampir maintenance YR3                                      | 3  | Sep-12 | Vampir upgraded   |
| 4.0     | <b>Operations</b>   |    |        |   |
| 4.1     | User Support  |    |        |   |
| 4.1.1   | Global research NOC network monitoring integration complete | 1  | Jan-10 | Service Desk monitoring of all FutureGrid networking components active      |
| 4.1.2.1 | IU KB entries created - Program Year 1                      | 1  | Sep-10 | 75 total KB entries available   |
| 4.1.2.2 | IU KB entries created - Program Year 2                      | 2  | Sep-11 | 150 total KB entries available  |
| 4.1.2.3 | IU KB entries created - Program Year 3                      | 3  | Sep-12 | 225 total KB entries available  |
| 4.1.2.4 | IU KB entries created - Program Year 4                      | 4  | Sep-13 | 300 total KB entries available  |
| 4.1.3.1 | Help Desk training on FutureGrid complete                   | 1  | Jan-10 | Help Desk ready for calls on FutureGrid functionality and support processes |
| 4.1.3.2 | Help Desk – PY2   | 2  | Sep-11 | Additional Help Desk training completed                                     |
| 4.1.3.3 | Help Desk – PY3   | 3  | Sep-12 | Additional Help Desk training completed                                     |
| 4.1.3.4 | Help Desk – PY4   | 4  | Sep-13 | Additional Help Desk training completed                                     |
| 5.0     | <b>Training, Education, and Outreach</b>                    |    |        |   |
| 5.1     | Conferences   |    |        |   |
| 5.1.1   | SC09  | 1  | Nov-09 | Initial demonstrations  |
| 5.1.2   | SC10  | 2  | Nov-10 | New FutureGrid capabilities available for demonstrations                    |
| 5.1.3   | SC11  | 3  | Nov-11 | New FutureGrid capabilities available for demonstrations                    |
| 5.1.4   | SC12  | 4  | Nov-12 | New FutureGrid capabilities available for demonstrations                    |
| 5.2     | Annual Surveys  |    |        |   |
| 5.2.1   | 2009  | 1  | Dec-09 | Feedback on what works well, what needs improvement, and enhancements       |
| 5.2.2   | 2010  | 2  | Dec-10 | Feedback on what works well, what needs improvement, and enhancements       |
| 5.2.3   | 2011  | 3  | Dec-11 | Feedback on what works well, what needs improvement, and enhancements       |
| 5.2.4   | 2012  | 4  | Dec-12 | Feedback on what works well, what needs improvement, and enhancements       |
| 5.3     | Coursework  |    |        |   |
| 5.3.1   | FutureGrid used in TACC classes                             | 2  | Oct-11 | Assessment of class and how well FutureGrid worked                          |
| 5.3.2   | New FutureGrid course at TACC completed                     | 4  | Sep-13 | Pre-packaged virtual machine images bundled with course material            |
| 6.0     | <b>Project Management</b>                                   |    |        |   |
| 6.1     | <b>Integrated Project Plans</b>                             |    |        |   |
| 6.1.1   | IPP Year 1 completed  | 1  | Oct-09 | Integrated Project Plan Y1 available on web site                            |
| 6.1.2   | IPP Year 2 completed  | 2  | May-10 | Integrated Project Plan Y2 available on web site                            |



FutureGrid: An Experimental, High-Performance Grid Test-Bed

| WBS      | Milestone  | YR | Finish | Metric  |
|----------|--|----|--------|---|
| 6.1.3    | IPP Year 3 completed                               | 3  | May-11 | Integrated Project Plan Y3 available on web site                    |
| 6.1.4    | IPP Year 4 completed                               | 4  | May-12 | Integrated Project Plan Y4 available on web site                    |
| 6.2      | <b>Status Reports</b>                              |    |        |   |
| 6.2.1    | Quarterly  |    |        |   |
| 6.2.1.1  | Q1 Y1 completed                                    | 1  | Mar-10 | Quarterly Status Report Q1 Y1 sent to NSF and available on web site |
| 6.2.1.2  | Q2 Y1 completed                                    | 1  | Jun-10 | Quarterly Status Report Q2 Y1 sent to NSF and available on web site |
| 6.2.1.3  | Q3 Y1 completed                                    | 1  | Oct-10 | Quarterly Status Report Q3 Y1 sent to NSF and available on web site |
| 6.2.1.4  | Q1 Y2 completed                                    | 2  | Mar-11 | Quarterly Status Report Q1 Y2 sent to NSF and available on web site |
| 6.2.1.5  | Q2 Y2 completed                                    | 2  | Jun-11 | Quarterly Status Report Q2 Y2 sent to NSF and available on web site |
| 6.2.1.6  | Q3 Y2 completed                                    | 2  | Oct-11 | Quarterly Status Report Q3 Y2 sent to NSF and available on web site |
| 6.2.1.7  | Q1 Y3 completed                                    | 3  | Mar-12 | Quarterly Status Report Q1 Y3 sent to NSF and available on web site |
| 6.2.1.8  | Q2 Y3 completed                                    | 3  | Jun-12 | Quarterly Status Report Q2 Y3 sent to NSF and available on web site |
| 6.2.1.9  | Q3 Y3 completed                                    | 3  | Oct-12 | Quarterly Status Report Q3 Y3 sent to NSF and available on web site |
| 6.2.1.10 | Q1 Y4 completed                                    | 4  | Mar-13 | Quarterly Status Report Q1 Y4 sent to NSF and available on web site |
| 6.2.1.11 | Q2 Y4 completed                                    | 4  | Jun-13 | Quarterly Status Report Q2 Y4 sent to NSF and available on web site |
| 6.2.1.12 | Q3 Y4 completed                                    | 4  | Sep-13 | Quarterly Status Report Q3 Y4 sent to NSF and available on web site |
| 6.2.2    | Annual   |    |        |   |
| 6.2.2.1  | Program Year 1 ( <i>October 2009 - Sept 2010</i> ) | 2  | Mar-11 | Program Year 1 Annual Report sent to NSF and available on web site  |
| 6.2.2.2  | Program Year 2 ( <i>October 2010 - Sept 2011</i> ) | 3  | Mar-12 | Program Year 2 Annual Report sent to NSF and available on web site  |
| 6.2.2.3  | Program Year 3 ( <i>October 2011 - Sept 2012</i> ) | 4  | Mar-13 | Program Year 3 Annual Report sent to NSF and available on web site  |
| 6.2.2.4  | Program Year 4 ( <i>October 2012 - Sept 2013</i> ) | 5  | Mar-14 | Program Year 4 Annual Report sent to NSF and available on web site  |
| 6.2.3    | "Lessons Learned" report completed                 | 5  | Apr-14 | "Lessons Learned" report published                                  |
| 6.3      | <b>Annual NSF Reviews</b>                          |    |        |   |
| 6.3.1    | Program Year 1 ( <i>October 2009 - Sept 2010</i> ) | 2  | Apr-11 | FutureGrid Annual Review at NSF                                     |
| 6.3.2    | Program Year 2 ( <i>October 2010 - Sept 2011</i> ) | 3  | Apr-12 | FutureGrid Annual Review at NSF                                     |
| 6.3.3    | Program Year 3 ( <i>October 2011 - Sept 2012</i> ) | 4  | Apr-13 | FutureGrid Annual Review at NSF                                     |
| 6.3.4    | Program Year 4 ( <i>October 2012 - Sept 2013</i> ) | 5  | Apr-14 | FutureGrid Annual Review at NSF                                     |
| 6.4      | <b>Annual FutureGrid Meeting</b>                   |    |        |   |
| 6.4.1    | Program Year 1                                     | 1  | Mar-10 | FutureGrid Annual Meeting at IU                                     |
| 6.4.2    | Program Year 2                                     | 2  | Mar-11 | FutureGrid Annual Meeting at IU                                     |
| 6.4.3    | Program Year 3                                     | 3  | Mar-12 | FutureGrid Annual Meeting at IU                                     |
| 6.4.4    | Program Year 4                                     | 4  | Mar-13 | FutureGrid Annual Meeting at IU                                     |
| 6.6      | <b>Project Web Site</b>                            |    |        |   |

FutureGrid: An Experimental, High-Performance Grid Test-Bed

| WBS   | Milestone             | YR | Finish | Metric  |
|-------|-----------------------|----|--------|---|
| 6.6.1 | Design completed      | 1  | Aug-09 | Final design approved   |
| 6.6.2 | Development completed | 1  | Sep-09 | Project web site developed and tested                                   |
| 6.6.3 | Web site deployed     | 1  | Oct-09 | Project web site used for on-going management and support of FutureGrid |

## Appendix B: WBS Dictionary

Below is the FutureGrid Work Breakdown Structure to at least project level 3. In some places where more detail seems appropriate, additional detail is provided.

| WBS   | Activity Name                     | Description  |
|-------|-----------------------------------|--|
| 1.0   | <b>Hardware</b>                   | This category encompasses all activities related to the procurement, installation, and implementation of computer resources at each FutureGrid site, including contractual acceptance benchmarks.  |
| 1.1   | <b>Dell 1152 core</b>             |  |
| 1.1.1 | Hardware configurations finalized |  |
| 1.1.2 | Vendor purchase orders finalized  |  |
| 1.1.3 | Site preparation                  | Tasks targeting any electrical, cooling, facility modifications or enhancements necessary for installation of the Dell 1152  |
| 1.1.4 | Hardware acquisition              | Tasks targeting the receipt and installation of the Dell 1152 and successfully connecting it to the FutureGrid network   |
| 1.1.5 | Commissioning                     | Tasks targeting contractual acceptance criteria for the Dell 1152, including the execution of benchmarks, installation and configuration of the FutureGrid software environment, complete systems testing activities, and a transition to operations                       |
| 1.2   | <b>IBM iDataPlex 1024 core</b>    |  |
| 1.2.1 | Hardware configurations finalized |  |
| 1.2.2 | Vendor purchase orders finalized  |  |
| 1.2.3 | Site preparation                  | Tasks targeting any electrical, cooling, facility modifications or enhancements necessary for installation of the IBM iDataPlex 1024 core  |
| 1.2.4 | Hardware acquisition              | Tasks targeting the receipt and installation of the IBM iDataPlex 1024 core and successfully connecting it to the FutureGrid network   |
| 1.2.5 | Commissioning                     | Tasks targeting contractual acceptance criteria for the IBM iDataPlex 1024 core, including the execution of benchmarks, installation and configuration of the FutureGrid software environment, complete systems testing activities, and a transition to operations support |
| 1.2.6 | Hardware refresh                  | Tasks targeting any requisite equipment refurbishing or necessary replacement of computer parts for the IBM iDataPlex 1024 core  |
| 1.3   | <b>IBM iDataPlex 672 core</b>     |  |
| 1.3.1 | Hardware configurations finalized |  |
| 1.3.2 | Vendor purchase orders finalized  |  |
| 1.3.3 | Site preparation                  | Tasks targeting any electrical, cooling, facility modifications or enhancements necessary for installation of the IBM iDataPlex 672 core   |
| 1.3.4 | Hardware acquisition              | Tasks targeting the receipt and installation of the IBM iDataPlex 672 core and successfully connecting it to the FutureGrid network  |
| 1.3.5 | Commissioning                     | Tasks targeting contractual acceptance criteria for the IBM iDataPlex 672 core, including the execution of benchmarks, installation and configuration of the FutureGrid software environment, complete systems testing activities, and a transition to operations          |

| WBS   | Activity Name                     | Description   |
|-------|-----------------------------------|---|
| 1.3.6 | Hardware refresh                  | Tasks targeting any requisite equipment refurbishing or necessary replacement of computer parts for the IBM iDataPlex 672 core  |
| 1.4   | IBM iDataPlex 256 core            |   |
| 1.4.1 | Hardware configurations finalized |   |
| 1.4.2 | Vendor purchase orders finalized  |   |
| 1.4.3 | Site preparation                  | Tasks targeting any electrical, cooling, facility modifications or enhancements necessary for installation of the IBM iDataPlex 256 core  |
| 1.4.4 | Hardware acquisition              | Tasks targeting the receipt and installation of the IBM iDataPlex 256 core and successfully connecting it to the FutureGrid network   |
| 1.4.5 | Commissioning                     | Tasks targeting contractual acceptance criteria for the IBM iDataPlex 256 core, including the execution of benchmarks, installation and configuration of the FutureGrid software environment, complete systems testing activities, and a transition to operations support |
| 1.4.6 | Hardware refresh                  | Tasks targeting any requisite equipment refurbishing or necessary replacement of computer parts for the IBM iDataPlex 256 core  |
| 1.5   | Cray XT5M 672 core                |   |
| 1.5.1 | Hardware configurations finalized |   |
| 1.5.2 | Vendor purchase orders finalized  |   |
| 1.5.3 | Site preparation                  | Tasks targeting any electrical, cooling, facility modifications or enhancements necessary for installation of the Cray XT5M 672 core  |
| 1.5.4 | Hardware acquisition              | Tasks targeting the receipt and installation of the Cray XT5M 672 core and successfully connecting it to the FutureGrid network   |
| 1.5.5 | Commissioning                     | Tasks targeting contractual acceptance criteria for the Cray XT5M 672 core, including the execution of benchmarks, installation and configuration of the FutureGrid software environment, complete systems testing activities, and a transition to operations             |
| 1.5.6 | Hardware refresh                  | Tasks targeting any requisite equipment refurbishing or necessary replacement of computer parts for the Cray XT5M 672 core  |
| 1.6   | Shared Memory cluster             |   |
| 1.6.1 | Hardware configurations finalized |   |
| 1.6.2 | Vendor purchase orders finalized  | Final PO to selected shared memory cluster vendor sent  |
| 1.6.3 | Site preparation                  | Tasks targeting any electrical, cooling, facility modifications or enhancements necessary for installation of the selected shared memory cluster  |
| 1.6.4 | Hardware acquisition              | Tasks targeting the receipt and installation of the selected shared memory cluster and successfully connecting it to the FutureGrid network   |
| 1.6.5 | Commissioning                     | Tasks targeting contractual acceptance criteria for the selected shared memory cluster, including the execution of benchmarks, installation and configuration of the FutureGrid software environment, complete systems testing activities, and a transition to operations |
| 1.7   | IBM iDataPlex 672 core            |   |

| WBS    | Activity Name                               | Description  |
|--------|---|--|
| 1.7.1  | IBM iDataPlex 672 core decommissioned at IU | Tasks targeting the removal of the IBM iDataPlex 672 core from Indiana University's data center for subsequent shipment to UCSD  |
| 1.7.2  | Site preparation                            | Tasks targeting any electrical, cooling, facility modifications or enhancements necessary for installation of the IBM iDataPlex 672 core machine being sent to UCSD  |
| 1.7.3  | Hardware acquisition                        | Tasks targeting the receipt and installation of the IBM iDataPlex 672 core at UCSD and successfully connecting it to the FutureGrid network  |
| 1.7.4  | Commissioning                               | Tasks targeting installation and configuration of the FutureGrid software environment, complete systems testing activities and a transition to operations  |
| 1.8    | Purdue "High Throughput" Cluster            |  |
| 1.9    | DataDirect Networks Storage Appliance       | DataDirect Networks S2A6620 120 TB Storage Appliance   |
| 1.9.1  | Vendor purchase orders finalized            | Final PO to DataDirect Networks sent   |
| 1.9.2  | Site preparation                            | Tasks targeting any electrical, cooling, facility modifications or enhancements necessary for installation of the DDN S2A6620 Storage Appliance at UC  |
| 1.9.3  | Hardware acquisition                        | Tasks targeting the receipt and installation of the DDN S2A6620 Storage Appliance at UC  |
| 1.9.4  | Commissioning                               | Tasks targeting contractual acceptance criteria for the DDN S2A6620 Storage Appliance, complete systems testing activities, and a transition to operations at UC   |
| 1.10   | SunFire Storage Servers                     | Two SunFire 4170 storage servers with each with Intel E5520 processors, 24GB of memory and 36TB of direct attached storage   |
| 1.10.1 | Vendor purchase orders finalized            | Final PO to Incentra sent  |
| 1.10.2 | Site preparation                            | Tasks targeting any electrical, cooling, facility modifications or enhancements necessary for installation of the SunFire Storage Servers at UCSD  |
| 1.10.3 | Hardware acquisition                        | Tasks targeting the receipt and installation of the SunFire Storage Servers at UCSD  |
| 1.10.4 | Commissioning                               | Tasks targeting contractual acceptance criteria for the SunFire Storage Servers, complete systems testing activities, and a transition to operations at UCSD   |
| 1.11   | HPSS  | Tasks related to the addition of volume and performance to the High performance Storage System (HPSS), including the procurement of additional tapes   |
| 2.0    | <b>Networks</b>                             | This category encompasses all activities related to network connectivity between all sites, including the procurement, installation, and implementation of network devices.  |
| 2.1    | Network contracts finalized                 | Tasks targeting the finalization of all network contracts with Starlight, Corporation for Network Education Initiatives in California (CENIC), National LambdaRail (NLR), Florida LambdaRail (FLR), AT&T, Spirent Communications, and Matrix Integration |
| 2.2    | Cisco 6509 Core Router                      | Cisco 6509 Core Router   |

| WBS     | Activity Name   | Description   |
|---------|---|---|
| 2.2.1   | Install/configure router                                  | Tasks targeting the receipt and installation of the Cisco 6509 Core Router at UC  |
| 2.3     | Spirent H10 XGEM Network Impairment emulator              | Spirent H10 XGEM Network Impairment emulator  |
| 2.3.1   | Install/configure network impairment simulator            | Tasks targeting the receipt and installation of the Spirent Network Impairment emulator at UC   |
| 2.4     | Connectivity to/from core router                          | Tasks targeting the successful connection to and from the Cisco Core Router by IU, UC, UF, UCSD, and TACC   |
| 2.5     | Perform network tests                                     | Tasks associated with the execution of all requisite networking tests on the entire FutureGrid network  |
| 2.6     | Provision connectivity to existing TeraGrid and Internet2 | Tasks associated with connecting Future Grid to the TeraGrid and Internet2  |
| 3.0     | <b>Software</b>   |   |
| 3.1     | AMIE  | Tasks targeting the connection of FutureGrid to the account allocation and management software used by the TeraGrid   |
| 3.2     | User Portal   | Tasks targeting the development and implementation of a web-based portal to provide a variety of functionality to both users and administrators. Example functionality includes requesting and managing resources for experiments, configuring resources, managing experiments, collaboration tools for user groups, documentation, and general monitoring of FutureGrid. |
| 3.2.1   | Initial version (Information Services)                    | Tasks targeting the development of the initial version of the portal, including its formal design, authentication and authorization, links to help information, and resource availability tracking  |
| 3.2.2   | Experiment Management Services                            | Tasks related to the delivery of services via portlets from the OGCE-based FutureGrid portal.   |
| 3.2.2.1 | Image Browser   | Tasks related to the inspection of information about images available for use in FutureGrid via the portal  |
| 3.2.2.2 | Experiment Browser  | Tasks related to the definition of experiments as resource, software, and experimental via the portal   |
| 3.2.2.3 | Software Configuration Browser                            | Tasks related to the specification of packages and configuration parameters for use in an experiment via the portal   |
| 3.2.2.4 | Monitoring/Instrumentation Browser                        | Tasks related to the examination of data gathered during experiments via the portal   |
| 3.2.2.5 | Scheduling, reservations                                  | Tasks related to the matching of researcher requests for test environments against availability via the portal  |
| 3.2.2.6 | Storage services  | Tasks related to the storage and retrieval of all software images and data relevant to a researcher's experiments   |
| 3.2.3   | View/manage user/group information                        | Tasks targeting the addition of user/group information management functionality in the portal   |
| 3.2.4   | Test harness access                                       | Tasks targeting access to the test harness via the portal   |
| 3.2.5   | Portal maintenance - PY2 H1                               | Tasks targeting bug fixes and those minor/approved enhancements to the portal in the first half of Program Year 2   |
| 3.2.6   | Portal maintenance - PY2 H2                               | Tasks targeting bug fixes and those minor/approved enhancements to the portal in the second half of Program Year 2  |

| WBS     | Activity Name  | Description   |
|---------|--|---|
| 3.2.7   | Portal maintenance – PY3 H1                                    | Tasks targeting bug fixes and those minor/approved enhancements to the portal in the first half of Program Year 3   |
| 3.2.8   | Portal maintenance – PY3 H2                                    | Tasks targeting bug fixes and those minor/approved enhancements to the portal in the second half of Program Year 3  |
| 3.2.9   | Portal maintenance – PY4 H1                                    | Tasks targeting bug fixes and those minor/approved enhancements to the portal in the first half of Program Year 4   |
| 3.2.10  | Portal maintenance – PY4 H2                                    | Tasks targeting bug fixes and those minor/approved enhancements to the portal in the second half of Program Year 4  |
| 3.3     | Pegasus  | Tasks targeting the open source software from USC Information Science Institute, providing implementation of experiment plans as workflow.  |
| 3.3.1   | Pegasus available on test-bed                                  | Tasks targeting the initial deployment of Pegasus on the FutureGrid network   |
| 3.3.2   | Pegasus documentation  | Tasks targeting the development and distribution of system and user documentation on Pegasus  |
| 3.3.3   | Immediate resource provisioning workflow                       | Tasks targeting the development and implementation of immediate resource provisioning workflows   |
| 3.3.4   | Time-sensitive resource provisioning workflow                  | Tasks targeting the development and implementation of time-sensitive resource provisioning workflows  |
| 3.3.5   | Workflow repository requirements                               | Tasks related to the gathering of requirements for the development of a workflow repository   |
| 3.3.6   | Pegasus tutorial   | Tasks targeting the development and distribution of an on-line tutorial on Pegasus  |
| 3.3.7   | End-to-end experiment management workflows                     | Tasks related to the development and implementation of new end-to-end workflows, from resource provisioning to injection of events available  |
| 3.3.8   | Workflow repository  | Tasks related to the development and implementation of the workflow repository  |
| 3.4     | Grid Benchmark Challenge                                       | Tasks targeting the development of a set of grid benchmarks to measure, characterize, and understand distributed application performance. These benchmarks will include a set of tightly coupled application grid benchmarks based on UTK's well-known HPC Challenge benchmarks and a set of loosely coupled application benchmarks based on real-world scientific workflow applications. UTK will define appropriate and relevant metrics for the performance, reliability, and variability of grid platforms and tightly coupled grid applications. These metrics will be deployed so that applications, architectures, and middleware implementations can evolve guided by sound engineering principles. |
| 3.5     | Inca   | Tasks targeting the user-level grid monitoring from UCSD as the standard monitoring tool for FutureGrid. Detects grid infrastructure problems by executing periodic, automated, user-level testing of grid software and services.   |
| 3.5.1   | Initial version  | Tasks targeting the initial deployment if Inca on the FutureGrid network  |
| 3.5.1.1 | Testing plan for monitoring FutureGrid functionality completed | Tasks related to the identification of what FutureGrid tests to deploy in Inca and the creation of a test plan to manage them   |

| WBS     | Activity Name   | Description  |
|---------|---|--|
| 3.5.1.3 | User documentation complete   | Tasks targeting the development and distribution of user documentation on Inca   |
| 3.5.1.5 | Support NSF required and optional benchmarks                        | Tasks related to deploying NSF benchmarks in Inca  |
| 3.5.2   | Add additional tests/benchmarks as new software is added or updated | Tasks related to upgrading Inca with new tests and benchmarks during PY2   |
| 3.5.3   | Add additional tests/benchmarks as new software is added or updated | Tasks related to upgrading Inca with new tests and benchmarks during PY3   |
| 3.5.4   | Add additional tests/benchmarks as new software is added or updated | Tasks related to upgrading Inca with new tests and benchmarks during PY4   |
| 3.6     | Nimbus  | Tasks targeting the open-source toolkit from UC that, once installed on a cluster, provides Infrastructure as a Service cloud to its client.                               |
| 3.6.1   | Nimbus deployment on UC and UF clusters                             | Tasks targeting the deployment of Nimbus on both the UC and UF IBM iDataPlex clusters  |
| 3.6.2   | Nimbus maintenance - PY2  | Tasks related to upgrading Nimbus with bug fixes and enhancements in PY2   |
| 3.6.3   | Nimbus maintenance - PY3  | Tasks related to upgrading Nimbus with bug fixes and enhancements in PY3   |
| 3.6.4   | Nimbus maintenance - PY4  | Tasks related to upgrading Nimbus with bug fixes and enhancements in PY4   |
| 3.7     | Actuating Services  | Tasks related to the Test-bed Management, Experiment Initiation, Collection of Experimental Data, Storage of Data for Later Retrieval and Use                              |
| 3.7.1   | Components  | The components of the Actuating Services are well-known, robust, open source software tools.   |
| 3.7.1.1 | Dagman  |  |
| 3.7.1.2 | Bcfg2   | Support for the bcfg2 service so that experiment workflows are automatically managed   |
| 3.7.1.3 | CondorG interface to MOAB/TORQUE                                    |  |
| 3.7.1.4 | MOAB/TORQUE interface to xCAT                                       |  |
| 3.7.2.1 | Xen   | Tasks targeting the instantiation of Xen, a virtual machine monitor  |
| 3.7.2.2 | Eucalyptus  | Tasks targeting the instantiation of Eucalyptus  |
| 3.7.2.3 | Nimbus  | Tasks targeting the instantiation of Nimbus  |
| 3.7.2.4 | VMWare  | Tasks targeting the instantiation of VMWare  |
| 3.7.2.5 | RPM-based Linux   | Tasks targeting the instantiation of an RMP-based Linux machine  |
| 3.7.2.6 | Windows 2008  | Tasks targeting the instantiation of a machine running Windows 2008  |
| 3.7.2.7 | Microsoft HPC Server  | Tasks targeting the instantiation of a machine running Microsoft HPC Server  |
| 3.8     | ViNe  | Virtual networking approach for grids from University of Florida, enabling symmetric connectivity among grid resources and allows existing applications to run unmodified. |




| WBS    | Activity Name                          | Description   |
|--------|--|---|
| 3.8.1  | ViNe routing software                  | Tasks related to integration of ViNe with Nimbus and other FutureGrid middleware  |
| 3.8.2  | ViNe management interfaces             | Tasks related to the specification of programmatic ViNe management APIs   |
| 3.8.3  | ViNe management services               | Tasks related to the development and initial deployment of programmatic ViNe management APIs  |
| 3.8.4  | ViNe routing and services improvements | Tasks related to upgrading ViNe for 1) monitoring and automatic recovery during network outages; 2) self-optimization of communication performance; and 3) end-to-end QoS   |
| 3.9    | Virtual appliance                      | Used to create unique, hands-on educational modules based on virtual appliance and social networking technologies from University of Florida. Easily boot up a prepackaged FutureGrid educational appliance on a user's own desktop and connect to a social network of other deployed FutureGrid appliances deployed over the network. This enables unique usage scenarios in education and training. |
| 3.9.1  | Initial version                        | Tasks targeting the initial version of the virtual appliance with Social VPN, including bindings to social networking back ends Open Social, Skype, and Gmail chat  |
| 3.9.2  | Education modules                      | Tasks related to the development of 3-4 education modules that utilize the virtual appliance  |
| 3.9.3  | Virtual appliance enhancements - PY3   | Tasks related to upgrading the virtual appliance software with bug fixes and enhancements in PY3  |
| 3.9.4  | Virtual appliance enhancements - PY4   | Tasks related to upgrading the virtual appliance software with bug fixes and enhancements in PY4  |
| 3.10   | Test Harness                           | Tasks related to new development from TACC that will allow FutureGrid users to efficiently execute one or more distributed experiments on configured machines. Examples of supported tasks include: scattering agents, programs, and files to machines; starting experiments; gathering experimental results during and after experiments; stopping experiments.                                      |
| 3.10.1 | Initial Version                        | Tasks related to the design, development, and implementation of the first version of the test harness, providing file transfers; start/stop agents, and a command-line interface  |
| 3.10.2 | Logging                                | Tasks targeting the merging of distributed logs into a unified experiment   |
| 3.10.3 | Web Interface                          | Tasks related to the design, development, and implementation of a web interface to the test harness for managing experiments  |
| 3.10.4 | Test harness maintenance - PY2 H1      | Tasks related to upgrading the test harness with bug fixes and enhancements in the 1st half of PY2  |
| 3.10.5 | Test harness maintenance - PY2 H2      | Tasks related to upgrading the test harness with bug fixes and enhancements in the 2nd half of PY2  |
| 3.10.6 | Test harness maintenance - PY3 H1      | Tasks related to upgrading the test harness with bug fixes and enhancements in the 1st half of PY3  |
| 3.10.7 | Test harness maintenance - PY3 H2      | Tasks related to upgrading the test harness with bug fixes and enhancements in the 2nd half of PY3  |

| WBS    | Activity Name   | Description  |
|--------|---|--|
| 3.10.8 | Test harness maintenance - PY4 H1                           | Tasks related to upgrading the test harness with bug fixes and enhancements in the 1st half of PY4   |
| 3.10.9 | Test harness maintenance - PY4 H2                           | Tasks related to upgrading the test harness with bug fixes and enhancements in the 2nd half of PY4   |
| 3.11   | Genesis II, UNICORE, and g-lite                             | Tasks related to the deployment of open source, standards-based grid platform (Genesis II) from University of Virginia designed to support both high-throughput computing and secure data sharing. |
| 3.11.1 | Acquire and train on UNICORE and g-lite                     | Tasks related to acquiring and learning the UNICORE and g-lite software  |
| 3.11.2 | Install UNICORE, and g-lite on local UV nodes               | Tasks related to the installation of UNICORE and g-lite on the UV IBM iDataPlex node   |
| 3.11.3 | Deploy UNICORE, and g-lite on FutureGrid nodes              | Tasks related to the deployment of UNICORE and g-lite on all FutureGrid nodes  |
| 3.11.4 | Deploy Genesis II on FutureGrid nodes                       | Tasks related to the deployment of Genesis II on FutureGrid nodes  |
| 3.11.5 | Deploy standard service endpoints for compliance testing    | Tasks related to the deployment of standard service endpoints for compliance testing   |
| 3.11.6 | Genesis II, UNICORE, and g-lite maintenance - PY2           | Tasks related to deploying bug fixes and enhancements to Genesis II, UNICORE, and g-lite for PY2   |
| 3.11.7 | Genesis II, UNICORE, and g-lite maintenance - PY3           | Tasks related to deploying bug fixes and enhancements to Genesis II, UNICORE, and g-lite for PY3   |
| 3.11.8 | Genesis II, UNICORE, and g-lite maintenance - PY4           | Tasks related to deploying bug fixes and enhancements to Genesis II, UNICORE, and g-lite for PY4   |
| 3.12   | Vampir  | Tasks related to the use of software from GWT-TUD GmbH that supports the analysis of applications performance in VM environments   |
| 3.12.1 | Deploy VampirServer on central server at IU                 | Tasks related to the deployment of VampirServer on central server at IU  |
| 3.12.2 | Deploy Vampir Trace on all FutureGrid nodes                 | Tasks related to the deployment of VampirTrace on all FutureGrid nodes   |
| 3.12.3 | Vampir maintenance - PY2                                    | Tasks related to upgrading Vampir with bug fixes and enhancements in PY2   |
| 3.12.4 | Vampir maintenance - PY3                                    | Tasks related to upgrading Vampir with bug fixes and enhancements in PY3   |
| 4.0    | <b>Operations</b>   |  |
| 4.1    | <b>User Support</b>   |  |
| 4.1.1  | Global research NOC network monitoring integration complete | Tasks related to the integration of network monitoring at the Global NOC for FutureGrid  |
| 4.1.2  | IU Knowledge Base   | Tasks related to the creation of FutureGrid entries into the Knowledge Base at IU  |
| 4.1.3  | Help Desk   | Tasks related to the central Help desk at IU becoming proficient in FutureGrid and servicing all Tier 1 problems   |
| 4.2    | <b>Computing Operations</b>                                 |  |
| 4.2.1  | Computer Operations - Program Year 1                        | System administration tasks associated with all FutureGrid clusters in PY1   |
| 4.2.2  | Computer Operations - Program Year 2                        | System administration tasks associated with all FutureGrid clusters in PY2   |


| WBS     | Activity Name   | Description  |
|---------|---|--|
| 4.2.3   | Computer Operations - Program Year 3  | System administration tasks associated with all FutureGrid clusters in PY3   |
| 4.2.4   | Computer Operations - Program Year 4  | System administration tasks associated with all FutureGrid clusters in PY4   |
| 4.3     | <b>Advanced User Support</b>  |  |
| 4.3.1   | Instantiating virtual clusters  | Tasks related to supporting FutureGrid users on how to instantiate a virtual cluster   |
| 4.3.2   | Configuring direct hardware requests  | Tasks related to supporting FutureGrid users on how to configure direct hardware requests  |
| 4.3.3   | Application installation and optimization (CPU and I/O) through profiling tools | Tasks related to learning FutureGrid profiling tools for supporting FutureGrid users on application installation and optimization                          |
| 5.0     | <b>Training, Education, and Outreach</b>  |  |
| 5.1     | <b>Conferences</b>  | Tasks related to the submission of papers, creation of demonstrations, BoFs, and participation on panels at conferences throughout the world on FutureGrid |
| 5.1.1   | SC09  | SC09, November 14-20, 2009   |
| 5.1.2   | SC10  | SC10, 2010   |
| 5.1.3   | SC11  | SC11, 2011   |
| 5.1.4   | SC12  | SC12, 2012   |
| 5.2     | <b>Annual Surveys</b>   | Tasks related to the creation, distribution, analysis, and publication of results from FutureGrid annual surveys   |
| 5.2.1   | Program Year 1  | Tasks related to the creation, distribution, analysis, and publication of results from FutureGrid PY1 annual survey  |
| 5.2.2   | Program Year 2  | Tasks related to the creation, distribution, analysis, and publication of results from FutureGrid PY2 annual survey  |
| 5.2.3   | Program Year 3  | Tasks related to the creation, distribution, analysis, and publication of results from FutureGrid PY3 annual survey  |
| 5.2.4   | Program Year 4  | Tasks related to the creation, distribution, analysis, and publication of results from FutureGrid PY4 annual survey  |
| 5.3     | <b>Coursework</b>   |  |
| 5.3.1   | FutureGrid tutorials  | Tasks related to IU preparation and publication of FutureGrid tutorials and other on-line materials  |
| 5.3.1.1 | FutureGrid tutorials - PY1  | Tasks related to IU preparation and publication of FutureGrid tutorials and other on-line materials in PY1   |
| 5.3.1.2 | FutureGrid tutorials – PY2  | Tasks related to IU preparation and publication of FutureGrid tutorials and other on-line materials in PY2   |
| 5.3.1.3 | FutureGrid tutorials – PY3  | Tasks related to IU preparation and publication of FutureGrid tutorials and other on-line materials in PY3   |
| 5.3.1.4 | FutureGrid tutorials – PY4  | Tasks related to IU preparation and publication of FutureGrid tutorials and other on-line materials in PY4   |
| 5.3.2   | Nimbus tutorials, on-line materials, etc.                                       | Tasks related to UC preparation and posting of Nimbus tutorials and other on-line materials over the course of the FutureGrid project                      |
| 5.3.3   | Social appliance tutorials, on-line materials, etc.                             | Tasks related to UF preparation and posting of social appliance tutorials and other on-line materials over the course of the FutureGrid project            |

| WBS     | Activity Name                                  | Description  |
|---------|--|--|
| 5.3.4   | Pegasus tutorials, on-line materials, etc.     | Tasks related to USC preparation and posting of Pegasus tutorials and other on-line materials over the course of the FutureGrid project  |
| 5.3.5   | TACC coursework                                | Tasks related to the use of FutureGrid in TACC classes and the development of pre-packaged virtual machine images to accompany coursework for other institutions be available as part of TACC course materials |
| 5.4     | <b>Outreach</b>                                |  |
| 5.4.1   | Open Grid Forum, EGEE, and Unicore             | Open Grid Forum, EGEE, and Unicore   |
| 5.4.2   | Alladin/Grid5K                                 | Alladin/Grid5K   |
| 5.4.3   | German D-Grid                                  | German D-Grid  |
| 5.4.4   | Minority Serving Institutions                  | Minority Serving Institutions  |
| 6.0     | <b>Project Management</b>                      |  |
| 6.1     | <b>Integrated Project Plans</b>                |  |
| 6.1.1   | Integrated project planning for Program Year 1 | Tasks related to the preparation of the Integrated Project Plan for FutureGrid PY1   |
| 6.1.2   | Integrated project planning for Program Year 2 | Tasks related to the preparation of the Integrated Project Plan for FutureGrid PY2   |
| 6.1.3   | Integrated project planning for Program Year 3 | Tasks related to the preparation of the Integrated Project Plan for FutureGrid PY3   |
| 6.1.4   | Integrated project planning for Program Year 4 | Tasks related to the preparation of the Integrated Project Plan for FutureGrid PY4   |
| 6.2     | <b>Status Reports</b>                          |  |
| 6.2.1   | Quarterly                                      | Tasks related to the creation and publication of all annual reports to the NSF   |
| 6.2.2   | Annual   | Tasks related to the creation and publication of all quarterly reports to the NSF  |
| 6.3     | <b>Annual NSF Reviews</b>                      | Tasks related to the preparation and publication of all materials required by NSF as part of their annual review, and the formal presentation to the NSF   |
| 6.4     | <b>Annual FutureGrid Meeting</b>               | Tasks related to the preparation, hosting, and documenting results of all FutureGrid annual meetings   |
| 6.5     | <b>Advisory Boards</b>                         |  |
| 6.5.1   | Science Advisory Board                         |  |
| 6.5.1.1 | Recruitment                                    | Tasks related to the initial recruitment for the FutureGrid Science Advisory Board   |
| 6.5.1.2 | Initial SAB meeting                            | Tasks related to the preparation, hosting, and documenting results of the first Science Advisory Board meeting   |
| 6.5.2   | User Advisory Committee                        |  |
| 6.5.2.1 | Recruitment                                    | Tasks related to the initial recruitment for the FutureGrid User Advisory Committee  |
| 6.5.2.2 | Initial UAC meeting                            | Tasks related to the preparation, hosting, and documenting results of the first User Advisory Committee meeting  |
| 6.6     | <b>Project Web Site</b>                        | Tasks related to the design, development, and implementation of the FutureGrid project web site  |

## Appendix C: FutureGrid Project Schedule




## FutureGrid: An Experimental, High-Performance Grid Test-Bed






## FutureGrid: An Experimental, High-Performance Grid Test-Bed

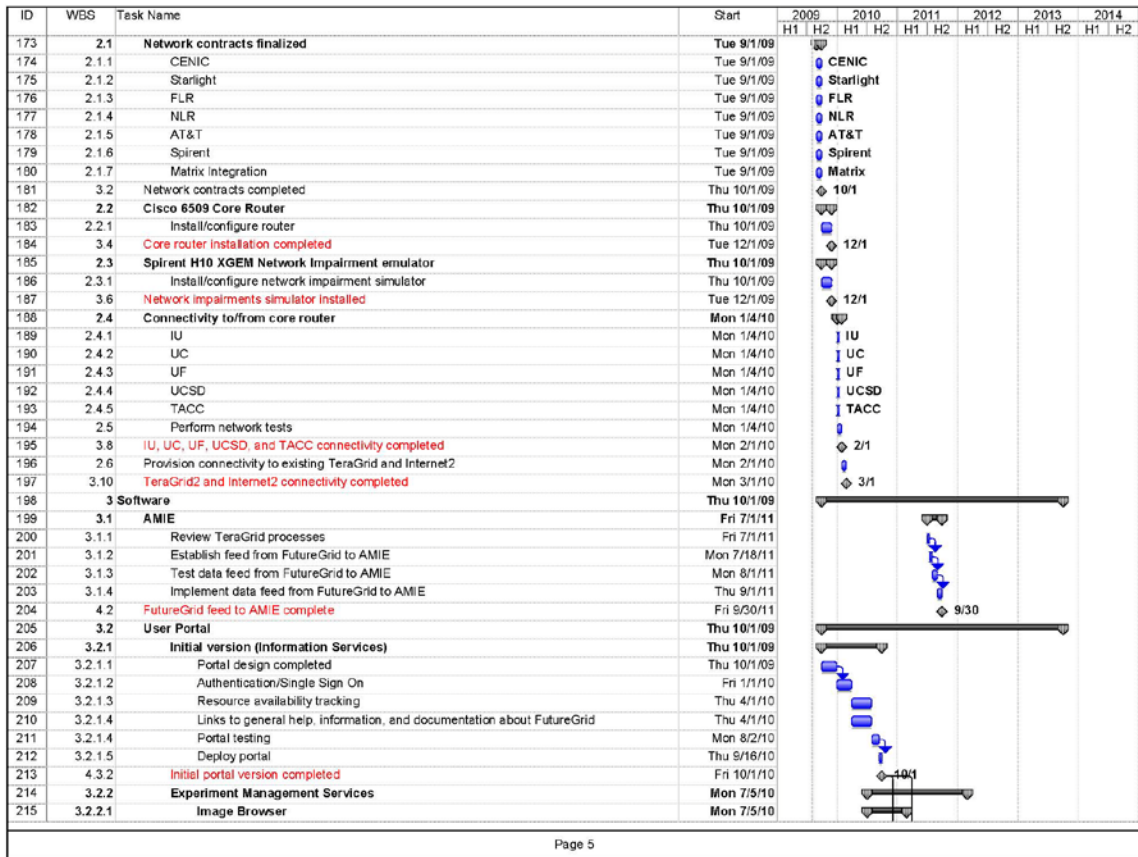


## FutureGrid: An Experimental, High-Performance Grid Test-Bed







## FutureGrid: An Experimental, High-Performance Grid Test-Bed




## FutureGrid: An Experimental, High-Performance Grid Test-Bed




# FutureGrid: An Experimental, High-Performance Grid Test-Bed




## FutureGrid: An Experimental, High-Performance Grid Test-Bed



## FutureGrid: An Experimental, High-Performance Grid Test-Bed




## FutureGrid: An Experimental, High-Performance Grid Test-Bed






## FutureGrid: An Experimental, High-Performance Grid Test-Bed




## FutureGrid: An Experimental, High-Performance Grid Test-Bed





## FutureGrid: An Experimental, High-Performance Grid Test-Bed



## FutureGrid: An Experimental, High-Performance Grid Test-Bed

| ID  | WBS     | Task Name                               | Start       | 2009 |    | 2010 |    | 2011 |    | 2012 |    | 2013 |    | 2014 |    |
|-----|---------|---|-------------|------|----|------|----|------|----|------|----|------|----|------|----|
|     |         |   |             | H1   | H2 | H1   | H2 | H1   | H2 | H1   | H2 | H1   | H2 | H1   | H2 |
| 560 | 6.4.4   | Program Year 4                          | Fri 3/1/13  |      |    |      |    |      |    |      |    |      |    |      |    |
| 561 | 8.4.9   | PY1 Annual FutureGrid Meeting completed | Mon 4/1/13  |      |    |      |    |      |    |      |    |      |    |      |    |
| 562 | 6.5     | Advisory Boards                         | Tue 9/1/09  |      |    |      |    |      |    |      |    |      |    |      |    |
| 563 | 6.5.1   | Science Advisory Board                  | Tue 9/1/09  |      |    |      |    |      |    |      |    |      |    |      |    |
| 564 | 6.5.1.1 | Recruitment                             | Tue 9/1/09  |      |    |      |    |      |    |      |    |      |    |      |    |
| 565 | 6.5.1.2 | Initial SAB meeting                     | Thu 10/1/09 |      |    |      |    |      |    |      |    |      |    |      |    |
| 566 | 6.5.2   | User Advisory Committee                 | Thu 10/1/09 |      |    |      |    |      |    |      |    |      |    |      |    |
| 567 | 6.5.2.1 | Recruitment                             | Thu 10/1/09 |      |    |      |    |      |    |      |    |      |    |      |    |
| 568 | 6.5.2.2 | Initial UAC meeting                     | Mon 11/2/09 |      |    |      |    |      |    |      |    |      |    |      |    |
| 569 | 6.5.3   | Executive Committee                     | Fri 10/2/09 |      |    |      |    |      |    |      |    |      |    |      |    |
| 570 | 6.6     | Project Web Site                        | Mon 8/3/09  |      |    |      |    |      |    |      |    |      |    |      |    |
| 571 | 6.6.1   | Design                                  | Mon 8/3/09  |      |    |      |    |      |    |      |    |      |    |      |    |
| 572 | 6.6.2   | Development                             | Tue 9/1/09  |      |    |      |    |      |    |      |    |      |    |      |    |
| 573 | 6.6.3   | Deployment                              | Thu 10/1/09 |      |    |      |    |      |    |      |    |      |    |      |    |
| 574 | 8.7     | Web site completed                      | Mon 10/5/09 |      |    |      |    |      |    |      |    |      |    |      |    |

*Appendix D: Projected Annual Cost by WBS*

| WBS #      | Description                                       | FY2010           | FY2011           | FY2012           | FY2012           | Total             |
|------------|---|------------------|------------------|------------------|------------------|-------------------|
| <b>1.0</b> | <b>Hardware</b>                                   | <b>2,617,243</b> | <b>715,050</b>   | <b>400,000</b>   | <b>85,000</b>    | <b>3,817,293</b>  |
| 1.1        | TACC Dell 1152 core                               | 501,563          |                  |                  |                  | 501,563           |
| 1.2.5.1    | IBM iDataPlex 1024 core acceptance test completed | 825,174          |                  |                  |                  | 825,174           |
| 1.3.5.1    | IBM iDataPlex 672 core acceptance test completed  | 513,346          |                  |                  |                  | 513,346           |
| 1.4.5.1    | IBM iDataPlex 256 core acceptance test completed  | 101,635          |                  |                  |                  | 101,635           |
| 1.5.5.1    | Cray XT5M 672 core acceptance test completed      | 470,000          |                  |                  |                  | 470,000           |
| 1.6.5.1    | Shared Memory cluster acceptance test completed   |                  | 630,050          |                  |                  | 630,050           |
| 1.9.4.1    | DataDirect Networks S2A6620 Storage Appliance     | 135,995          |                  |                  |                  | 135,995           |
| 1.10.4.1   | Sun X4540 96 TB X4540 Storage Server              | 69,530           |                  |                  |                  | 69,530            |
| 1.11       | HPSS  | —                | 10,000           | 150,000          | 10,000           | 170,000           |
| 1.2.6      | Hardware Refresh (also 1.3.6, 1.4.6, and 1.5.6)   |                  | 75,000           | 250,000          | 75,000           | 400,000           |
| <b>2.0</b> | <b>Networks</b>                                   | <b>450,605</b>   | <b>163,869</b>   | <b>163,869</b>   | <b>163,869</b>   | <b>942,212</b>    |
| 2.2        | Cisco 6509 Core Router                            | 114,407          |                  |                  |                  | 114,407           |
| 2.3        | Spirent H10 XGEM Network Impairment emulator      | 77,020           | 11,554           | 11,554           | 11,554           | 111,682           |
| 2.4.1      | IU connected to core router                       | 211,655          | 122,315          | 122,315          | 122,315          | 578,600           |
| 2.4.2      | UC connected to core router                       | 30,000           | 6,000            | 6,000            | 6,000            | 48,000            |
| 2.4.3      | UF connected to core router                       | 5,523            | —                | —                | —                | 5,523             |
| 2.4.4      | UCSD connected to core router                     | 12,000           | 24,000           | 24,000           | 24,000           | 84,000            |
| <b>3.0</b> | <b>Software</b>                                   | <b>1,215,575</b> | <b>1,405,393</b> | <b>1,442,039</b> | <b>1,360,645</b> | <b>5,423,652</b>  |
| 3.1,3.7    | AMIE, FutureGrid Software Environment             | 522,690          | 535,288          | 551,572          | 500,788          | 2,110,338         |
| 3.2        | User Portal                                       | 121,021          | 124,651          | 128,391          | 132,243          | 506,306           |
| 3.3        | Pegasus   | 95,566           | 93,925           | 92,187           | 38,789           | 320,467           |
| 3.4        | Grid Benchmark Challenge                          | —                | 98,016           | 101,922          | 105,983          | 305,921           |
| 3.5        | Inca  | 150,792          | 155,316          | 159,975          | 164,774          | 630,857           |
| 3.6        | Nimbus  | 131,040          | 134,971          | 139,020          | 143,191          | 548,222           |
| 3.8,3.9    | ViNe, SocialVPN                                   | 12,194           | 13,500           | 13,905           | 14,322           | 53,921            |
| 3.10       | Test Harness                                      | 105,983          | 109,162          | 112,438          | 115,810          | 443,393           |
| 3.11       | Genesis II, UNICORE, and g-lite                   | 76,289           | 78,194           | 80,259           | 82,375           | 317,117           |
| 3.12       | Vampir  | —                | 62,370           | 62,370           | 62,370           | 187,110           |
| <b>4.0</b> | <b>Operations</b>                                 | <b>693,447</b>   | <b>803,692</b>   | <b>829,580</b>   | <b>808,860</b>   | <b>3,162,814</b>  |
| 4.1        | User Support                                      | 131,611          | 152,575          | 157,152          | 161,866          | 603,204           |
|            | IU  | 131,611          | 152,575          | 157,152          | 161,866          | 603,204           |
| 4.2        | Computing Operations                              | 561,836          | 651,117          | 672,428          | 646,994          | 2,559,610         |
|            | IU  | 222,092          | 228,755          | 235,618          | 242,686          | 929,151           |
|            | UC  | 134,784          | 138,828          | 142,992          | 147,282          | 563,886           |
|            | UCSD  | 18,661           | 19,220           | 24,025           | 7,208            | 69,114            |
|            | UF  | 122,391          | 126,063          | 129,845          | 133,741          | 512,040           |
|            | USC   | 63,908           | 63,393           | 62,889           | 36,855           | 227,045           |
|            | UTK   | 0                | 1,470            | 1,470            | 1,366            | 4,306             |
|            | TACC  |                  | 73,388           | 75,589           | 77,857           | 226,833           |
|            | UV  | 6,607            | 6,740            | 6,875            | 7,013            | 27,235            |
| <b>5.0</b> | <b>Training, Education, and Outreach</b>          | <b>149,115</b>   | <b>185,101</b>   | <b>183,730</b>   | <b>184,787</b>   | <b>702,733</b>    |
|            | IU  | 56,501           | 59,382           | 58,843           | 60,745           | 235,471           |
|            | UC  | 37,244           | 38,361           | 39,512           | 40,698           | 155,815           |
|            | UCSD  | 9,778            | 10,072           | 10,374           | 10,685           | 40,909            |
|            | UF  | 0                | 163              | 668              | 1,184            | 2,015             |
|            | USC   | 19,840           | 20,677           | 21,543           | 22,460           | 84,520            |
|            | UTK   | 0                | 30,226           | 26,088           | 21,916           | 78,230            |
|            | TACC  | 15,601           | 16,069           | 16,551           | 17,047           | 65,268            |
|            | UV  | 10,151           | 10,151           | 10,151           | 10,052           | 40,505            |
| <b>6.0</b> | <b>Project Management</b>                         | <b>\$436,759</b> | <b>\$375,631</b> | <b>\$374,519</b> | <b>\$463,979</b> | <b>1,650,888</b>  |
|            | IU  | 304,159          | 223,098          | 216,894          | 301,079          | 1,045,230         |
|            | UC  | 37,244           | 38,361           | 39,512           | 40,698           | 155,815           |
|            | UCSD  | 9,778            | 10,072           | 10,374           | 10,685           | 40,909            |
|            | UF  | 27,542           | 30,490           | 31,403           | 32,345           | 121,780           |
|            | USC   | 26,218           | 27,538           | 28,912           | 30,354           | 113,022           |
|            | UTK   | 0                | 13,461           | 14,000           | 14,560           | 42,021            |
|            | TACC  | 15,601           | 16,069           | 16,551           | 17,047           | 65,268            |
|            | UV  | 16,217           | 16,542           | 16,873           | 17,211           | 66,843            |
|            | <b>Grand Total WBS</b>                            |                  |                  |                  |                  | <b>15,699,592</b> |

*Appendix D.1: Projected Annual Cost by WBS NSF-IU Breakdown*

| WBS Number | Description                       | NSF Funded   | IU Funded   | Total        |
|------------|-----------------------------------|--------------|-------------|--------------|
| 1.0        | Hardware                          | \$1,835,883  | \$1,981,410 | \$3,817,293  |
| 2.0        | Networks                          | \$137,523    | \$804,689   | \$942,212    |
| 3.0        | Software                          | \$4,603,074  | \$820,578   | \$5,423,652  |
| 4.0        | Operations                        | \$1,371,873  | \$1,532,355 | \$2,904,228  |
| 5.0        | Training, Education, and Outreach | \$835,315    | \$126,004   | \$961,319    |
| 6.0        | Project Management                | \$1,316,332  | \$364,825   | \$1,681,157  |
|            | Grand Total WBS 1.0 - 6.0         | \$10,100,000 | \$5,629,861 | \$15,729,861 |

## Appendix E. FutureGrid **DRAFT** Staffing Plan

| Personnel                              | Role                            | FTE                           |                               |                               |                               |
|--|---------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|
|  |                                 | Y1                            | Y2                            | Y3                            | Y4                            |
| <i>(IU match given in parentheses)</i> |                                 |                               |                               |                               |                               |
| <i>Indiana University</i>              |                                 |                               |                               |                               |                               |
| <i>Total NSF + match given in bold</i> |                                 |                               |                               |                               |                               |
| Geoffrey Fox                           | PI                              | 0.20<br>(0.05)<br><b>0.25</b> | 0.00<br>(0.25)<br><b>0.25</b> | 0.00<br>(0.25)<br><b>0.25</b> | 0.20<br>(0.05)<br><b>0.25</b> |
| Craig Stewart                          | Senior Personnel                | 0.20<br><b>0.20</b>           | 0.03<br>(0.17)<br><b>0.20</b> | 0.00<br>(0.20)<br><b>0.20</b> | 0.19<br><b>0.19</b>           |
| Gregor von Laszewski                   | Software Architect              | 0.80                          | 0.80                          | 0.80                          | 0.80                          |
| David Hancock                          | Senior Personnel                | 0.25<br><b>0.25</b>           | 0.07<br>(0.18)<br><b>0.25</b> | 0.00<br>(0.25)<br><b>0.25</b> | 0.20<br>(0.05)<br><b>0.25</b> |
| Ray Sheppard                           | Advanced support                | (1.00)                        | (1.00)                        | (1.00)                        | (1.00)                        |
| Marlon Pierce                          | Gateway architect               | (0.25)                        | (0.25)                        | (0.25)                        | (0.25)                        |
| Jonathon Bolte                         | Online support manager          | 0.00<br>(0.20)<br><b>0.20</b> | 0.00<br>(0.20)<br><b>0.20</b> | 0.00<br>(0.20)<br><b>0.20</b> | 0.00<br>(0.20)<br><b>0.20</b> |
| Gary Miksik                            | Project manager                 | 0.50<br><b>0.50</b>           | 0.08<br>(0.42)<br><b>0.50</b> | 0.00<br>(0.50)<br><b>0.50</b> | 0.00<br>(0.50)<br><b>0.50</b> |
| Richard Knepper                        | Site lead                       | (0.50)                        | (0.50)                        | (0.50)                        | (0.50)                        |
| Joseph Rinkovsky                       | System admin                    | (1.00)                        | (1.00)                        | (1.00)                        | (1.00)                        |
| Siddharth Maini                        | Gateway developer               | (1.00)                        | (1.00)                        | (1.00)                        | (1.00)                        |
| Brent Sweeny                           | Network engineer                | (0.50)                        | (0.50)                        | (0.50)                        | (0.50)                        |
| (Fraction of 3 person, 24/7 team)      | GRNOC support                   | (0.25)                        | (0.50)                        | (0.50)                        | (0.50)                        |
| TBN                                    | Programmer/analyst              | 1.00<br><b>1.00</b>           | 1.00<br><b>1.00</b>           | 1.00<br><b>1.00</b>           | 0.50<br>(0.50)<br><b>1.00</b> |
| TBN                                    | Programmer/analyst              | 0.75<br>(0.25)<br><b>1.00</b> | 0.75<br>(0.25)<br><b>1.00</b> | 0.75<br>(0.25)<br><b>1.00</b> | 0.00<br>(1.00)<br><b>1.00</b> |
| <i>University of Tennessee</i>         |                                 |                               |                               |                               |                               |
| Jack Dongarra                          | Senior Personnel                | 0.00                          | 0.00                          | 0.00                          | 0.00                          |
| Piotr Luszczek                         | Senior Personnel                | 0.00                          | 0.10                          | 0.10                          | 0.10                          |
| TBN                                    | Post-doc                        | 0.00                          | 0.67                          | 0.67                          | 0.67                          |
| TBN                                    | Graduate Student                | 0.00                          | 1.00                          | 1.00                          | 1.00                          |
| <i>Texas Advanced Computing Center</i> |                                 |                               |                               |                               |                               |
| Warren Smith                           | Co-PI                           | 0.15                          | 0.15                          | 0.15                          | 0.15                          |
| Maytal Dahan                           | Senior Personnel                | 0.25                          | 0.25                          | 0.25                          | 0.25                          |
| Patrick Hurley                         | Programmer/Analyst              | 0.50                          | 0.50                          | 0.50                          | 0.50                          |
| David Carver                           | System Administrator/Programmer | 0.00                          | 0.50                          | 0.50                          | 0.50                          |

FutureGrid: An Experimental, High-Performance Grid Test-Bed

| Personnel                          | Role                  | FTE  |      |      |      |
|------------------------------------|-----------------------|------|------|------|------|
|                                    |                       | Y1   | Y2   | Y3   | Y4   |
| University of California-San Diego |                       |      |      |      |      |
| Shava Smallen                      | Senior Personnel      | 0.20 | 0.20 | 0.20 | 0.20 |
| Philip Papadopoulos                | Senior Personnel      | 0.03 | 0.03 | 0.03 | 0.03 |
| TBN                                | Programmer/Analyst    | 1.00 | 1.00 | 1.00 | 1.00 |
| TBN                                | System administrator  | 0.11 | 0.11 | 0.11 | 0.03 |
| University of Chicago              |                       |      |      |      |      |
| Katarzyna Keahey                   | Co-PI                 | 0.40 | 0.40 | 0.40 | 0.40 |
| Ian Foster                         | Senior Personnel      | 0.00 | 0.00 | 0.00 | 0.00 |
| TBN                                | Programmer            | 1.00 | 1.00 | 1.00 | 1.00 |
| Ti Leggett                         | System administration | 1.00 | 1.00 | 1.00 | 1.00 |
| University of Florida              |                       |      |      |      |      |
| Jose Fortes                        | Co-PI                 | 0.11 | 0.11 | 0.11 | 0.11 |
| Renato Figueiredo                  | Senior Personnel      | 0.11 | 0.11 | 0.11 | 0.11 |
| TBN                                | System administrator  | 0.88 | 0.88 | 0.88 | 0.88 |
| University of Southern California  |                       |      |      |      |      |
| Ewa Deelman                        | Senior Personnel      | 0.10 | 0.10 | 0.10 | 0.10 |
| TBN                                | Programmer/Analyst    | 0.50 | 0.50 | 0.50 | 0.25 |
| University of Virginia             |                       |      |      |      |      |
| Andrew Grimshaw                    | Co-PI                 | 0.05 | 0.05 | 0.05 | 0.05 |
| TBN                                | Graduate Student      | 1.00 | 1.00 | 1.00 | 1.00 |
| TBN                                | Graduate Student      | 0.75 | 0.75 | 0.75 | 0.75 |

## *Appendix F. IU Notice to Subcontract Recipients*

This important notice has been disseminated to all FutureGrid participating entities, who understand that compliance with the terms of this Important Notice constitute a key element of the Interface Agreement. Compliance is a condition for participation in FutureGrid.



# Important Notice

**Subject: Subrecipients**

**No. 04-1**

**Date: January 26, 2004**

(Note: This Important Notice is being sent to Fiscal Officers, Chairpersons, Deans and Chancellors. Please forward to others who have a need to know.)

This Important Notice is issued to outline the subrecipient process at Indiana University and the responsibilities for subrecipient monitoring. Subrecipient agreements are sometimes referred to as subcontracts or consortium agreements. Federal OMB Circular A-133, "Audits of States, Local Governments, and Non-Profit Organizations", establishes audit requirements for federal and federal pass through funds received at Indiana University and other institutions of higher education. Under a prime federal award, Indiana University may desire work to be completed by an outside entity. Section §\_\_\_.210 of A-133 entitled "Subrecipient and Vendor Determinations" gives guidance in assessing if a subrecipient relationship exists. **The Indiana University practice is to extend similar requirements and definitions to non-federal subawards.**

**Definition of a Subrecipient:** A non-Federal entity that expends Federal awards from a pass-through entity to carry out a Federal program. A subrecipient may also be a recipient of other Federal awards directly from a Federal awarding agency.

Characteristics of a subrecipient:

- Receiving entity determines who is eligible to receive what Federal financial assistance;
- Has its performance measured against whether the objectives of the Federal program are met;
- Has responsibility for programmatic decision making;
- Has responsibility for adherence to applicable Federal programs compliance requirements; and
- Uses the Federal funds to carry out a program of the organization as compared to providing goods or services for a program of the pass-through entity.

**Definition of a Vendor:** A dealer, distributor, merchant or other seller providing goods or services that are required for the conduct of a Federal program. These goods or services may be for an organization's own use or for the use of beneficiaries of the Federal program.

Characteristics of a vendor:

- Provides the goods and services within normal business operations;
- Provides similar goods or services to many different purchasers;
- Operates in a competitive environment;
- Provides goods or services that are ancillary to the operation of the Federal program; and
- Is not subject to the compliance requirements of the Federal program.

Realizing that there may be unusual circumstances or exceptions to the above, you are encouraged to work with your research office in making the determination of whether a subrecipient or vendor relationship exists. Vendor agreements are handled by the Purchasing Department upon the initiation of a requisition by the department. Vendor agreements may be for consulting services, contractual agreements, or other fee for service arrangements and will follow normal procurement laws and regulations.

To comply with the requirements of OMB Circular A-133 for federal projects and to provide sound fiscal stewardship on all sponsored projects, Indiana University is responsible for monitoring subrecipients. The table below details the roles and responsibilities of subrecipient monitoring and subaward administration at our institution.

## SUBRECIPIENT MONITORING

| ROLE  | RESPONSIBILITY  |
|---|---|
| Determine if a subrecipient relationship exists.  | When submitting a proposal:   |
| IU Project Director,  |   |
| Department,   |   |
| The research office for your campus:  |   |
| IUPUI – Research and Sponsored Programs   |   |
| IUB and Regionals – Sponsored Research Services   |   |
| Collect the following information from the proposed subrecipient:   |   |
| Scope of work   |   |
| Budget and budget justification   |   |
| Institutional authorization   |   |
| Copy of indirect cost (F&A) rate agreement  | IU Project Director before routing proposal to the IU research office |
| Review subrecipient's budget to ensure that the funding agency's expense guidelines are followed. Work with proposed subrecipient to correct budget problems and to ensure that the correct indirect cost rate is proposed. | The research office for your campus:                                  |
| IUPUI – Research and Sponsored Programs   |   |
| IUB and Regionals – Sponsored Research Services   |   |



|   |  |
|---|--|
| Issue Subawards (also may be called Subcontracts or Consortium Agreements), negotiate terms, issue amendments, provide IU authorizing signature.  | IUPUI awards – Research and Sponsored Programs   |
| IUB and regional campus awards – Contract and Grant Administration  |  |
| Review subrecipient invoices to ensure that the funding agency's cost policies are followed and that expenditures fall within the dollar amount and time period of the agreement.   | IU Project Director, Department Fiscal Officer, Account Manager or Delegate  |
| Review subrecipient invoices to ensure that the appropriate program milestones are being met relative to the rate of expenditures.  | IU Project Director, Department Fiscal Officer, Account Manager or Delegate  |
| Collect program/technical reports as required by the subaward agreement and the prime funding agreement.  | IU Project Director, Department Fiscal Officer, Account Manager or Delegate  |
| Collect cost share verification (if required) and report subaward expenditures to prime sponsor.  | Contract and Grant Administration  |
| Ensure that the correct subcontract object code and indirect cost rate (F&A) have been applied to the IU account.   | Departmental Fiscal Officer or Account Manager, Contract and Grant Administration  |
| Collect A-133 audit reports or audited financial reports from all subrecipients of federal funds received through IU.   | Contract and Grant Administration  |
| Determine if there are disallowances.   | Contract and Grant Administration  |
| Collect refunds from subrecipient if disallowances are made.  | Department Fiscal Officer or Account Manager with assistance from Contract and Grant Administration  |
| Cover overdrafts caused by disallowances.   | Department   |
| Additional monitoring may be required if problems are found in the normal review steps listed above. Monitoring techniques may include but are not limited to: review of indirect cost (F&A) rate agreements, review of fringe benefit rates, desk audits of expenditures, and site visits. | Determined by Contract and Grant Administration in consultation with all of the above parties at the time of instituting additional monitoring procedures. |

### Additional References:

- OMB Circular A-133 <http://www.whitehouse.gov/omb/circulars/a133/a133.html>
- Important Notice 89-14, Subcontract Agreements with Small Organizations [http://www.ovpra.indiana.edu/cg/imp\\_notice/89-14.asp](http://www.ovpra.indiana.edu/cg/imp_notice/89-14.asp)
- FMS, Contract and Grant Administration, A-133 <http://www.ovpra.indiana.edu/cg/a133.asp>
- Research Gateway <http://www.research.indiana.edu/>